WHAT MAKES WRONGFUL DISCRIMINATION WRONG? BIASES, PREFERENCES, STEREOTYPES, AND PROXIES

LARRY ALEXANDER†

TABLE OF CONTENTS

INTRODUCTION	151
I. BACKGROUND ASSUMPTIONS	154
II. DISCRIMINATION AND PREFERENCES	157
A. <i>Preferences for and Against Particular Kinds of People</i> ..	158
1. Categorical Preferences for People: The Problem of Biases	158
2. Preferences for Particular Types of People as Reflections of Role Ideals	163
3. Personal Aversions and Attractions to Particular Types of People	165
B. <i>Transitional Cases Between Preferences For and Against Particular Types of People and Preferences for Specific Goods and Services: Proxies and Reactions</i>	167
1. Discrimination Based on Proxy Traits and Stereotypes	167
2. Reaction Qualifications	173
C. <i>Preferences for Goods, Services, and Personal Relations That Entail Preferences for Particular Kinds of People</i> ..	176
1. Preferences for Goods, Services, and Personal Relations Tainted by Associations with Biases, Stereotypes, and Ideologies	178
a. <i>Generally</i>	178
b. <i>The Special Case of Unconscious Bias</i>	179
2. Preferences for Goods, Services, and Relations That Differentially Affect Various Groups	183

† Professor of Law, University of San Diego. My debts to others are numerous for whatever virtues this article possesses. I would like to thank the law faculties at Boston University, IIT, Chicago-Kent, S.M.U., University of Pennsylvania, University of San Diego, and University of Toronto for inviting me to present drafts of the article at their workshops. In addition, I would like to thank the following individuals who gave me detailed criticism and suggestions: Elaine Alexander, Bob Bone, Kevin Cole, Richard Epstein, John Garvey, Kent Greenawalt, Michael Harper, Gail Heriot, Harry Hutchison, Richard Kay, Evan Lee, Michael Levin, Bill Marshall, Gay Morgan, Marty Redish, Maimon Schwarzschild, Emily Sherwin, Kristine Strachan, Alan Wertheimer, Chris Wonnell, and Fred Zacharias. Finally, I would like to thank my research assistants, Barry Ponticello and David Rankin, for their help.

a.	<i>Disparate Impact and Present Social Consequences</i>	185
b.	<i>Disparate Impact, Present Social Consequences, and Past Wrongs</i>	187
c.	<i>The Special(?) Case of Preferences for Traditional Ways of Doing Things</i>	188
III.	WHEN IS ACTING ON CERTAIN PREFERENCES MORALLY WRONG AND WHAT SHOULD BE DONE ABOUT IT?	190
A.	<i>Intrinsically Immoral Preferences</i>	192
B.	<i>Intrinsically Benign and Only Contingently Immoral Preferences</i>	194
1.	Effects That Bear on the Morality of Intrinsically Benign Preferences	194
a.	<i>Effects Caused by the Discrimination</i>	194
b.	<i>Effects on the Discriminator of Forgoing the Discrimination</i>	199
C.	<i>Immutability of Traits</i>	200
D.	<i>Immoral Discrimination and Moral Rights</i>	201
1.	The Moral Rights of the Discriminator	201
2.	The Moral Rights of Victims of Immoral Discrimination	202
3.	Immoral Discrimination and Moral Rights: A Summary	202
E.	<i>The Law and Immoral Discrimination</i>	203
F.	<i>Three Current Issues: Comparable Worth, Disparate Impact, and Affirmative Action</i>	208
1.	The Comparable Worth Issue	209
2.	The Disparate Impact Issue	212
3.	The Private Affirmative Action Issue	217
	CONCLUSION	218

INTRODUCTION

All of us well-socialized Westerners know that discrimination against other human beings is wrong. Yet we also realize, if we think about it at all, that we discriminate against others routinely and inevitably. We all know it is wrong to refuse to hire women as truck drivers, to refuse to let blacks practice law, to bar Moslems from basketball teams, or to refuse to sit next to Rastafarians at lunch counters. At the same time, we also know it is not wrong to refuse to hire the blind as truck drivers, to refuse to admit those who flunk the bar exam to the practice of law, to bar short, slow, uncoordinated persons from the basketball team, or to refuse to sit next to people who haven't bathed recently.

What explains and justifies the distinctions we make between discrimination that is wrongful and discrimination that is not? I argue in this article that answering this question is much more difficult than most people assume. Indeed, despite the fact that the morality and legality of discrimination have been at the forefront of academic inquiry for about three decades, the fundamental question about discrimination—What distinguishes wrongful discrimination from permissible forms of discrimination?—has seldom been addressed, much less answered.

Of course, I do not deny that some answers have been offered. Many people, at least when first asked, respond that basing discrimination on immutable traits such as race or gender is what makes discrimination wrong. The implausibility of this answer is exposed by the many instances where discrimination based on immutable traits is not regarded as wrong (for example, refusing to hire the blind as truck drivers), and by those instances of wrongful discrimination involving mutable characteristics (for example, barring Moslems from the basketball team). Similarly, labeling a trait on which discrimination is based as "irrelevant" begs the question of what makes the trait irrelevant. After all, the trait is relevant to the discriminator, whose purposes in choosing can be quite varied and complex.¹

¹ See JAN NARVESON, *THE LIBERTARIAN IDEA* 316-18 (1988):

Further conundrums loom as we reflect further. What constitutes a "job"? This is important for the following reason. If we are to try to spell out the ideas of "discrimination" in terms of "morally irrelevant," and in turn to make relevance turn on the job to be done, with its appropriate

criteria of competence, then the concept of wrongful discrimination would seem to be determined by the type of job it is—by the *job description*, as we shall call it (meaning, not just what is put in the advertisements, but rather the set of objectives that the holder of that job is, as such, to be pursuing). The idea is that it is wrong to hire persons for reasons other than those related to performance on the job—meaning, then, the job *as described*. But who describes the job? That is to say, whose job is it?

If the job is in the “private sector” of the economy, it would seem that this description is determined by the *owners* of the firm, normally through their higher-level employees, the managers. Is there any limit to the way in which they can *frame* a job description? For example, can they say “Wanted: Secretary/mistress: Successful applicants will be chosen on the basis of probable performance in the office and in bed. Here’s a picture of your boss-to-be: . . .”? Or “Wanted: black stenographer”? Or what about “Professor of Logic—definite advantage if applicant is black”?

On one view of the matter, they could do this. In the case of the mistress/secretary, they likely wouldn’t get a very good response to their ads, though (and perhaps the ones who did reply would be rejected as “unsuitable”—insufficiently sexually attractive, say). And then we may say that those who apply to ads that *didn’t* include the extra bit about bed performance and then are turned down for reasons actually having to do with that have indeed been unjustly treated, for the job has been *misrepresented*. They have been wasting their time applying, for the job they thought they were applying for isn’t the one they were being “looked over” for! (Suppose the job description said: “secretary/coffeemaker: in addition to usual secretarial skills, successful applicant is to make and serve coffee to senior staff, etc.” The inclusion of this not unusual supplement to a secretarial ad would, I think, solve a lot of problems, for many people would apply for such a job. Would those who would apply have a legitimate complaint? Is there fixed somewhere in the heavens an Essence of Secretary that clearly specifies that secretaries are *not* to make coffee for those whose secretaries they are?)

Similarly, if the claim about injustice is based on this consideration, then employers who want to discriminate on the basis of race, for instance, could also do so so long as their ads were properly worded. But the Ontario Civil Rights Commission would not allow this, and neither would the current version of the Canadian Charter of Rights and Freedoms, nor the current interpretation of the U.S. Constitution by American judges. Should these documents forbid such things?

There is a difficulty in doing so, since sometimes the forbidden values of [a characteristic possessed by one job candidate but not another] *will* be relevant to the job. For a chorus line or a model for women’s underwear, the employer does want a woman; for other purposes, a black person may be just what is needed (cf. the Black Muslim Church of America); and so on. On what grounds does the government decide whether this kind of distinction is “discrimination” or not? What if the Catholic Church discriminates against women for the priesthood? (There was recently a to-do in the papers concerning a girl who wanted to function as an acolyte, this being contrary to Catholic tradition. Did she have a case?) How about when the Black Muslim Church discriminates against white persons? Or the E[c]uadorian Friendship Society discriminates against non-E[c]uadorian applicants for janitorial or secretarial positions; or . . . These are all jobs

Other suggested answers point to the historical and contemporary statuses of various groups as the key to understanding what is "wrong" about wrongful discrimination. I believe these answers are closer to the mark, though in need of considerable refinement. Are all novel forms of discrimination—say, against blue-eyed persons—unproblematic merely because they do not target a group that historically has been the subject of widespread discrimination? Or given that the group of persons with low IQs and few academic achievements also suffers from very low socioeconomic status in contemporary society, is discrimination against such persons in the job market now wrongful?

In what follows I am going to approach the question of what makes discrimination wrongful by examining discrimination as an expression of various types of preferences. Part I briefly sets forth the framework that I am assuming in assessing the morality of discrimination. Part II examines various types of preferences and the discrimination to which they give rise. Part III attempts to formulate the results of the previous section insofar as they reveal the factors central to the wrongfulness of wrongful discrimination. Perhaps unsurprisingly, the results suggest that the line between wrongful and acceptable discrimination is, in most cases, difficult to locate with precision because it is historically and culturally variable. This line is historically and culturally variable because it is, in most cases, a function of consequentialist considerations rather than deontological norms. That is, in most cases, discrimination, when it is wrongful, is contingently but not intrinsically so.

A word about the form this inquiry takes. The reader will no doubt find herself from time to time becoming increasingly impatient with the multitude of distinctions drawn—distinctions among types of discriminatory preferences and among the contexts in which they occur—and with the rather elaborate taxonomy I construct out of those distinctions. I understand that impatience, especially because the subject is so emotionally charged. Yet, for better or worse, those distinctions are my message. Discrimination is not one thing, but many. Failure to recognize this point results in intellectual and moral confusion as well as bad policy.

Furthermore, the reader should not expect a sustained, rigorous philosophical argument. I offer no full-blown normative theory as

that *could* be done by women, or whites, or non-E[c]cuadorians, and so on.

a backdrop for my discussion of discrimination, nor any metaethical position. What the reader will find instead is a somewhat messy blend of deontological and consequentialist considerations brought to bear on a variety of forms and contexts of discrimination. The "philosophy" in the article, if it can be called that, is of a very rough, street-level, colloquial style. I am more interested in sorting things out than in wrapping them up, and what rigor there is lies in that sorting.

One final point: although there are no good systematic treatments of the morality of discrimination, I cite below five previous works of more limited scope that have been enormously helpful and influential in my thinking about the subject.² Any serious inquiry into the morality of discrimination should begin with them.

I. BACKGROUND ASSUMPTIONS

Many discussions of discrimination suffer from failure to be clear about the context in which discrimination is to be assessed. Discrimination by government officials may be morally as well as legally different from discrimination by private parties. Discrimination in the workplace may be morally as well as legally different from discrimination in one's choice of friends. And discrimination that exacerbates an unjust distribution of wealth may be morally different from discrimination that occurs within an otherwise just society.

Six assumptions provide the framework for my discussion of discrimination. First, I deal exclusively with discrimination by private parties, not discrimination by government officials. Although I believe that the morality of private discrimination bears heavily on the morality of public discrimination, I do not intend to discuss the latter in any way.

Second, I assume that in a just society there will be an area of liberty in which private people are permitted to express their preferences with respect to their intimate companions, their associates, their employees and employers, the salaries they pay and

² Those works are: GORDON W. ALLPORT, *NATURE OF PREJUDICE* (1954); HENRI TAJFEL, *HUMAN GROUPS AND SOCIAL CATEGORIES* (1981); Paul Brest, *Foreword: In Defense of the Antidiscrimination Principle*, 90 HARV. L. REV. 1 (1976); Richard A. Wasserstrom, *Racism, Sexism, and Preferential Treatment: An Approach to the Topics*, 24 UCLA L. REV. 581 (1977); Alan Wertheimer, *Jobs, Qualifications, and Preferences*, 94 ETHICS 99 (1983).

the salaries they work for, and the goods and services they consume. In other words, I assume as background that a just society is realized in a liberal democratic society with a sizeable free market sector and a rich realm of personal liberty and privacy. What I have to say will thus have much less relevance in a socialist society and none in a totalitarian one. Therefore, if someone believes that there can be no social justice in any society with a large domain of private freedom of choice, including a substantial free market, that person will find my project entirely misguided and diversionary. Why worry about the morality of a private choice to avoid black or female employees when justice would not countenance any employer freedom to choose employees?

The third assumption I make is perhaps the most crucial because it focuses on a point that engenders considerable confusion in discussions of discrimination. I assume as background the existence of a just society, one in which each person is guaranteed those political and civil rights and the minimum standard of living that justice requires. I am not going to argue for any particular theory of justice, nor will I claim that any particular society fully satisfies the correct theory of justice. I am only going to assume that, for purposes of assessing the morality of private discrimination, every individual receives the minimum rights and resources justice requires.

This third assumption will seem unusually strong to some readers. I make it, however, to avoid what I perceive is a constant source of confusion in the literature on discrimination. Our reactions to instances of discrimination are often colored by our sense that many victims of such discrimination have generally received a "raw deal" at the hands of society and face grinding poverty, a lack of education, and the concomitant bleak prospects for any kind of self-fulfilling, fully human existence. If we believe that a just society would not permit such human degradation, then we will also view private discrimination that reinforces or worsens such conditions as unjust. What I wish to ask, however, is whether there is anything morally distinctive about some kinds of discrimination apart from their contribution to social conditions that are unjust on independent grounds, or does achievement of a just society eliminate all potential moral criticisms of private discriminatory choices? In short, for purposes of moral assessment, I want to

isolate discriminatory choices from all the other wrongmaking factors with which they are frequently associated.³

A corollary of the previous two, my fourth assumption is that one can have a moral right to do what is morally wrong. The only dissent from this position among mainstream Western philosophers comes from act-consequentialists. For them, assessing an act of private discrimination is no different from assessing any other act: no act is intrinsically wrong, and any act may be wrong depending on its consequences.⁴ Putting aside the act-consequentialists, most philosophers agree that having the moral liberty to do X does not mean that doing X is either morally correct or free from moral criticism. Surely the libertarians, who argue for the widest realm of liberties, generally concede that exercise of those liberties can be subject to moral evaluation and criticism.⁵ Even more conventional liberals, who countenance a good deal more regulation of private choices, distinguish between having a moral right to choose and exercising that choice in a morally correct way.⁶ Most discussions of rights are attempts to fix those boundaries, and acts that fall within those boundaries represent virtually uncharted territory from the standpoint of moral analysis. This article is in part an attempt to begin mapping that domain.⁷

³ Another assumption, too closely related to the third assumption to warrant independent discussion, is that social justice is ordinarily, though not always, best promoted through macro policies addressed to wealth distribution, access to education, and reparations for denials of just entitlements, rather than through micro policies that target private choices.

An alternative to bracketing the issue of social justice in this manner would be to count among the morally relevant effects of private discrimination its aggravation—or, in some cases, its amelioration—of unjust conditions. I believe that this alternative would be much messier analytically than the bracketing alternative, although it may be the case that my preference for dealing with social justice through macro policies affects my perception of the relative analytical merits of these alternatives.

⁴ See, e.g., JOSEPH RAZ, *THE MORALITY OF FREEDOM* 268-69 (1986) (noting that strict consequentialists believe that no act is intrinsically wrong).

⁵ See, e.g., ROBERT NOZICK, *ANARCHY, STATE, AND UTOPIA* 31-32 (1974) (distinguishing between political philosophy, with its focus on rights, and moral philosophy, which is broader).

⁶ See, e.g., DAVID HAMLIN, *THE NAZI/SKOKIE CONFLICT* 50, 123 (1980) (acknowledging that Nazi doctrine is "monstrous and evil," yet supporting Nazis' right to demonstrate in the presence of Jewish counterdemonstrators); Alan Wertheimer, *Two Questions About Surrogacy and Exploitation*, 21 *PHIL. & PUB. AFF.* 211, 239 (1984) (stating that "[i]t is commonplace that the realm of . . . rights . . . is only part of our moral landscape").

⁷ My analysis will necessarily be different from a straight law and economics approach since the latter takes all preferences as givens—as exogenous factors—which are not themselves subject to a law and economics critique. My analysis speaks to our

In my fifth assumption, I part company with the more extreme libertarians. I assume that when an exercise of liberty is seriously morally wrong, there is reason to doubt that moral liberty extends to such exercise. I assume therefore that although there is a realm of moral liberty that covers most of our acts of discrimination, including morally wrong acts of discrimination, some seriously morally wrong acts of discrimination do not fall within this realm of moral liberty. Thus, determining that an act of discrimination is morally wrong gives us a reason, though not a conclusive one, to think that the act may not be within the realm of moral liberty. If the act is both morally wrong and outside the realm of moral liberty, we have a reason, though again not a conclusive one, to prohibit it by law and attach sanctions to its exercise.

My sixth and final assumption follows from the others. What I have to say about the morality of various types of discrimination will be relevant to, though not conclusive of, various legal concerns. Identifying a type of discrimination as morally wrong provides some reason for prohibiting it legally. A moral analysis of discrimination, therefore, might inform the interpretation of both statutory and constitutional law and should inform proposals for or against legal change. Although I am not primarily engaging in legal analysis, my inquiry is surely of major importance to the law.

II. DISCRIMINATION AND PREFERENCES

We discriminate against certain people and in favor of others because the satisfaction of our preferences leads us to do so. As they relate to discrimination, these preferences break down into two main divisions: preferences for and against certain people and preferences for various goods and services. Both kinds of preferences lead inevitably to discrimination. But discrimination that flows from preferences for and against people raises issues that are distinct from those raised by discrimination flowing from preferences for various goods and services.

The reader should keep in mind throughout the following discussion of types of preferences that determining which of these ideal types underlies any particular act of discrimination in the real world will often prove quite difficult or impossible. Many of these ideal types blend into one another and produce borderline examples that cannot be classified without controversy. Moreover,

foundational moral entitlements and is logically prior to efficiency considerations.

many real world preferences are arguably mixtures of the ideal types I describe. Nonetheless, the following purified taxonomy should guide any moral inquiry in the always messy and ambiguous empirical world.

A. Preferences for and Against Particular Kinds of People

When it comes to choosing our employees or employers, our colleagues, our friends, our spouses, or our regular golf groups, we are selective. Only the most unusual person makes these selections in a purely random fashion. The more virtuous among us might base their choices on morally valued qualities of character. Thus, they might seek associations with the wise, the just, the compassionate, the brave, and the loyal. Most of us look, however, for traits in our associates in addition to the moral virtues, traits that are less morally freighted, such as a sense of humor, a similarity of outlook, and physical attractiveness. And sometimes we look for traits that are more problematic, such as membership in a certain race, ethnic group, or religion.

Preferences for people, and the contexts in which those preferences arise, can be morally evaluated based on what they reveal about the moral qualities of those who have the preferences. Thus, a person who prefers a member of her own race as a spouse, but who is happy to work and play alongside members of other races, presents a different moral case from a person who prefers members of her own race in all contexts. The latter preference is a categorical racial preference, whereas the former is a contextual racial preference that is, for most people, much less morally troublesome.

1. Categorical Preferences for People: The Problem of Biases

Sometimes discrimination reflects the discriminator's biases for and against certain types of people. Biases, or differential moral concern, tend to be reflected in categorical preferences for and against certain types of people, that is, preferences that hold regardless of context. Thus, a Nazi who regards Aryans as more worthy of moral concern than others, and Jews as less worthy, will usually manifest these biases by preferring associations with Aryans and dispreferring associations with Jews in all contexts, from marriage to employment.

The Nazis' biases were both intrinsically morally wrong and profoundly devastating in their effects, effects that need not and cannot be recounted here. Their biases were intrinsically morally

wrong because Jews are clearly not of lesser moral worth than Aryans. When a person is judged incorrectly to be of lesser moral worth and is treated accordingly, that treatment is morally wrong regardless of the gravity of its effects. It represents a failure to show the moral respect due the recipient, a failure which is by itself sufficient to be judged immoral.

Not all biases are like those of the Nazis, however, even if we set effects aside. First, some biases are not uncontroversially intrinsically wrong. Many believe that all persons deserve equal concern and respect,⁸ but many others believe that the morally virtuous deserve more concern and respect than the morally vicious.⁹ For this group, a bias in favor of the virtuous and against the vicious is not only morally permissible but morally required.

It would be tempting to say that, except for biases in favor of the morally virtuous and against the morally vicious, all biases are morally illegitimate bases for preferences and the acts that express them. Surely this is the case with historically important biases, such as biases against members of particular races, genders, nationalities, or religions. No plausible theory regards blacks, females, the Irish, or the Jews as meriting less moral concern than whites, males, Germans, or Protestants merely by virtue of their status as a member of the former group. Of the multitude of everyday types of discrimination, based on aptitudes, skills, physical characteristics, and personality traits, no one seriously entertains the thought that those not preferred are less morally worthy than those preferred.

Yet it overstates the case to say moral virtue and vice are the only plausible bases for differential moral concern. Some people maintain that it is morally permissible, laudatory, or even obligatory to feel more concern for the welfare of kin, tribe, community, or nation than for the welfare of others. We do not simply tolerate ethnic, local, and national pride: we celebrate it. Many of us view as morally deficient rather than as properly unbiased one who shows no more concern for his own child than for that of another.¹⁰

The moral analysis of the realm of bias—differential moral concern—is messy, but some tentative conclusions can be drawn from common moral intuitions and reflection thereon. First, except

⁸ See, e.g., RONALD DWORKIN, *TAKING RIGHTS SERIOUSLY* 180-83 (1977); JOHN RAWLS, *A THEORY OF JUSTICE* 504-12 (1971).

⁹ See, e.g., GEORGE SHER, *DESERT* 142-49 (1987).

¹⁰ See Bernard Williams, *Persons, Character and Morality*, in *MORAL LUCK* 1 (Bernard Williams ed., 1981).

(controversially) for the traits identifying the morally virtuous and vicious, no other traits differentiate those who merit greater and lesser moral concern than others. Along all other axes, people are moral equals.

Second, personal commitments, relations, and identifications morally permit and may require particular persons to have greater moral concern for some than for others, even if the preferred individuals merit no greater moral concern from people in general because they possess no greater moral worth than others. My family and my neighbors are morally no more worthy and deserving of concern than others' families or neighbors, but they are certainly more deserving of *my* concern.

Third, morally favoring a small group and (relatively) morally disfavoring the rest of humanity has a different moral quality from morally disfavoring a small group. This is so primarily because the disfavoring of a small group is less likely to be the logical corollary of positive personal commitments and ties to others than it is to be the manifestation of an ideology that proclaims erroneously that members of the small group are morally unworthy. Put differently, my ties to the Alexanders do not require me to believe that the Joneses are morally inferior. Any "ties" I feel towards the white race, however, to the extent they produced anti-black bias, would most likely be based on an ideology of black moral inferiority.

The significance of this distinction between morally favoring large groups and morally favoring small groups should not be overstated. When the "small" groups approach the size of ethnic groups or nations, it is quite likely that group favoritism has as its corollary widespread belief in the moral inferiority of nonmembers. A belief that one's small ethnic group is a morally chosen people will quite naturally entail a concomitant belief in outsiders' moral inferiority.

Moreover, one reason why it is difficult to disentangle special concern for "one's own" from beliefs in others' moral inferiority is because in humankind's primitive past, these two attitudes were inextricably linked. When we roamed the earth in small kinship groups, and every tribe but one's own was a deadly enemy, love of one's kin and hatred and/or suspicion of everyone else were traits necessary for survival. We may be conditioned by our primitive past so that cosmopolitan moral views are much easier to attain intellectually than emotionally.¹¹ Without the emotional prop, a

¹¹ See Christopher T. Wonnell, *Circumventing Racism: Confronting the Problem of the*

belief in the equal moral worth of those outside one's tribe may be fragile. Leaving aside differential moral concern for the especially virtuous and vicious, one conclusion to be drawn from the discussion of biases to this point is that 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. On the other hand, because they are not based on judgments of differential moral worth, "biases" that reflect real personal attachments to family and locale are not intrinsically morally wrong.

There are two further points about biases that bear on whether one should have a moral right to act on them, even if doing so is morally wrong. First, biases that are intrinsically morally wrong because they reflect incorrect judgments of differential moral worth are usually not central to personal identity. To the extent they are fully conscious, they should be eradicable relatively easily through moral education, at least for those who have plausible sources of self-esteem besides believing that another group is morally inferior to theirs. The belief in blacks' intrinsic moral inferiority was probably not a deep-seated aspect of whites' personalities even in the period of its dominance, which explains the ideology's rather quick disappearance (in historical perspective), except as an isolated phenomenon associated with fringe groups.

Those who think the previous paragraph's relegation of anti-black bias to history's dustbin reflects willful blindness, naiveté, or malice should remember that I am dealing with bias, a judgment that those with a certain trait are morally less worthy than others merely by virtue of possessing that trait. I am not dealing with stereotypes and proxies, judgments that persons with a certain trait are quite likely to possess other traits that are proper bases for attributing differential moral worth. A person who believes, not that being black *per se* makes one morally less worthy, but that being black statistically correlates quite highly with being a criminal—a trait that more plausibly reflects lower moral worth—is not biased in the same sense that I am discussing in this section. However, inaccurate stereotypes are frequently the product of biases which have been disavowed intellectually but which still govern emotionally and create "tastes" for erroneous beliefs. Stereotypes and the proxy judgments based upon stereotypes are much more

Affirmative Action Ideology, 1989 B.Y.U. L. REV. 95, 106. But see TAJFEL, *supra* note 2, at 129-30 (expressing skepticism about socio-biological, as opposed to cognitive, explanations of prejudice).

difficult to eradicate than conscious biases because they frequently reflect accurate or at least plausible factual judgments and an understandable aversion to risk, and because when they do not, they are most likely the product of biases that are not fully conscious.¹²

I feel much less confident in what I have said about the relative lack of depth of conscious biases and their relatively easy eradication through moral education when we turn from bias against a particular group to biases that reflect a relatively small group's sense of its own moral superiority. Perhaps the latter biases are relatively deep features of the group's members' personalities. Nonetheless, I tend to doubt it. For remember that here we are dealing with judgments of differential moral worth. We are not dealing with personal loyalties and commitments, where differential moral concern is detached from judgments of differential moral worth. Nor are we dealing with stereotypes and proxies. Tribal loyalties and commitments and stereotypical judgments can account, I think, for most instances of what might otherwise appear to be judgments of one's own group's moral superiority.

The second point about biases is that they can have devastating social consequences. The more people within a society who are biased against another social group, the worse the consequences will be, because bias, as I have defined it, will show up in all contexts for choices, from the most intimate to the most public. Moreover, to the extent the biased group is generally better placed on the socioeconomic ladder than the group that is the target of bias, the social consequences are likely to be worse.

These harmful social consequences are likely to be both psychological and material. Members of the disfavored group will feel either stigmatized and inferior or insulted and angry. Their sense of well-being will in either case be reduced. Some may be moved to violence. Moreover, discrimination against them in the job market, a natural consequence of bias, will tend to reduce their aspirations, their energy, and consequently their productivity, to their detriment and to the detriment, in material terms at least, of the biased group as well.

These harmful social consequences depend upon the number of people among the biased and the disfavored groups and their relative socioeconomic status. If only a handful of people are biased, few if any harmful psychological or material effects will be

¹² For a full discussion of unconscious biases, stereotypes, and proxies, see *infra* Part II.C.1.b.

produced. If the disfavored group is socially and economically dominant, it is unlikely either to feel psychological harm or to have its productivity affected. Moreover, widespread bias against a disfavored group low on the socioeconomic ladder occasionally produces a sense of superiority and a redoubling of effort by that group rather than psychological injury or reduced motivation. Thus, the social effects of bias are contingent on a number of factors and will not be uniform amongst all societies, groups, and historical eras.

Where harmful social effects will ensue from bias, given the numbers and group characteristics, there is probably a case for legally prohibiting biased choices in certain realms otherwise left to private choice, particularly the economic realm. Apart from those directed at the exceptionally morally vicious/virtuous or those reflecting small group loyalties, biased choices are intrinsically morally wrong. There is therefore less reason to believe there is a moral right to make biased choices when they produce harmful consequences, even within a framework that meets the minimum standards of justice. Although attempting to extirpate all such immoral biased choices through law would be too costly and could violate moral rights, as in the case of prohibiting a choice of spouse based on bias, the adverse consequences of prohibiting biased economic choices by law might be worth the benefits of eliminating the various costs of biased economic choices. Of course, in practice it will be very difficult to distinguish immorally biased economic choices from choices that are not immorally biased and not intrinsically wrong, such as choices to favor one's family or friends or ethnic group members, or choices based on negative stereotypes. Thus, it is best that we survey the other forms of discrimination before reaching any conclusions, however tentative, about social policy.

2. Preferences for Particular Types of People as Reflections of Role Ideals

There is another category of preferences for and against types of people distinct from, but easily confused with, bias on the one hand, and stereotyping on the other. Some hold, *as a moral ideal*, that people with a particular trait should perform certain tasks and occupy certain social roles. The ideal cannot be based on a judgment of differential moral worth, or it would be an instance of

bias.¹³ Nor can the ideal be based on a presumed high statistical correlation between the defining trait and some other trait that is relevant to the task in question, or it would be an instance of stereotype.¹⁴ Moreover, the moral ideal on which the preferences in question are based must be distinguished from other ideals that produce preferences for particular types of people for particular social roles and occupations, including aesthetic ideals (basketball should be played only by people who are graceful or handsome as well as skilled), and technical ideals (no one should practice law without a good background in the liberal arts).

Perhaps the category of role preferences based on moral ideals is a null set. Arguably, however, some people believe as a matter of moral ideal that women and men should perform distinct social roles. Such people do not believe that men and women have differential moral worth. Nor do they believe that all women are technically better suited for certain roles than all men, or that all men are technically better suited for other roles than all women. Rather, they believe that it is immoral for women to perform certain roles and for men to perform certain roles. Although few people believe gender matters morally for every social role, some believe gender matters morally in child rearing and in sexuality. Thus, many Americans view heterosexuality as the only morally proper form of sexuality, and that the role of sex partner is gender-specific as a matter of morality and not just as a matter of taste.¹⁵ Arguably, these widespread beliefs about the morality of gender roles are in many cases based not on biases or on stereotypes but on moral ideals.

It is very difficult to reach any firm conclusions about discriminatory preferences based on moral ideals regarding roles. On the one hand, these moral ideals often are offered by their proponents without argument as self-evident or as religiously revealed truths, and it is hard to know what arguments for such ideals would look like. Those arguments for morally required roles which do emerge in debates over gender usually take the form of arguments about technical competence (women are better suited for child rearing), or look suspiciously like aesthetic appeals. They do not take the

¹³ See discussion *supra* Part II.A.1.

¹⁴ See discussion *infra* Part II.B.1.

¹⁵ See, e.g., *Bowers v. Hardwick*, 478 U.S. 186, 196 (1986) (holding that "majority sentiments about the morality of homosexuality" are an adequate basis for criminalizing homosexual sodomy).

form of moral arguments for gender-specific roles.¹⁶ On the other hand, I surely do not want to rule out the possibility of moral ideals regarding roles. All moral arguments begin with unargued-for premises.

Preferences based on moral ideals, if widespread, will be experienced as oppressive by all individuals who reject the moral ideals and their own role assignments thereunder. Moreover, they will be seen as oppressive by those who reject the ideals, even if they themselves are not consigned to unwanted roles, because of the effects on others who mistakenly accept the ideals and role assignments that prevent them from realizing their full potential for flourishing. If the ideals are mistaken, then the society in which they are widespread will suffer the costs of failed self-realization and lost productivity, even if that society does not perceive these as costs, or perceives them as costs that are morally required.

3. Personal Aversions and Attractions to Particular Types of People

Many people have aversions or attractions to particular types of people, either categorically or in specific contexts. These aversions and attractions are not biases because they are not based on judgments of differential moral worth, though these aversions and attractions may be rooted in ancient biases that now operate subconsciously. Nor are these aversions and attractions based on moral ideals or on stereotypes—those who have them do not believe that the unwanted (or wanted) associations are morally forbidden (or required), or that the target group likely has some other trait that is straightforwardly relevant to the association—though these aversions and attractions also may be rooted in moral ideals or in stereotypes that now operate subconsciously. Finally, I am not here concerned with aversions and attractions that are merely the corollaries of small group ties.

¹⁶ These statements reveal my inability to grasp the sense of those “natural law” moral arguments favored by, among others, the Catholic Church, which transform biological functions into moral imperatives (e.g., sex is morally permissible only for the purpose of procreation because that is its biological function). See Paul Bromberg, *Abortion and the Morality of Nurturance*, 21 CAN. J. PHIL. 513, 521-24 (1991) (placing the abortion debate in the framework of moral attitudes towards nurturance, one of which is the conservative argument that a woman’s primary moral duty is to have and nurture children).

Thus, some men may feel quite uncomfortable transacting corporate business with women without consciously believing that women are morally inferior to men, that women (morally) do not belong in corporate boardrooms, or that women are technically unqualified to handle corporate matters. Some women may feel quite uncomfortable with a male gynecologist without believing that males cannot or should not be gynecologists. These aversions and attractions are probably not central to the identities of those who have them, but often they may be rather deep-seated psychologically. Some may be traceable to biological hardwiring. Others are most likely the products of biases, ideals, and stereotypes that have become buried in the subconscious: they will tend, if widespread, to reinforce conscious biases, ideals, and stereotypes, and to be experienced by their victims in the same way the victims experience conscious biases, ideals, and stereotypes.¹⁷ If they are uniform, so that the same groups tend to be preferred and dispreferred, and widespread, and if they relegate the dispreferred to less desirable positions in the socioeconomic hierarchy, their social effects may be quite devastating. If they are idiosyncratic and variable, uncommon, or context-specific—"I'm uncomfortable around Italians in my private club but not at work"—rather than categorical—"I prefer to avoid Jews in all contexts"—and do not disprefer the already relatively disadvantaged, their adverse social effects may be relatively minimal.

Aversions and attractions based on race, ethnicity, or gender are only part of this category of preferences for and against certain types of people. More common perhaps are aversions and attractions based on aesthetics—how others look, sound, or even smell. Physical appearance is a frequent basis for preferences, either categorically ("I want to be around handsome people in all contexts") or in particular contexts ("I hire only good looking dancers for my troupe"). Although aversions and attractions based on physical attractiveness are common, they usually neither derive from nor reinforce biases, ideals, or stereotypes. Nor are these preferences so strong and widespread that those of us who are not stunningly

¹⁷ I attend to the moral status of aversions and attractions based on unconscious biases, a most difficult issue, below. See discussion *infra* Part II.C.1.b; see also Charles R. Lawrence, III, *The Id, the Ego, and Equal Protection: Reckoning with Unconscious Racism*, 39 STAN. L. REV. 317, 328-44 (1987) (explaining the origins and effects of unconscious racism).

beautiful are seriously affected by them, unless, of course, we pursue those careers that place a premium on looks.¹⁸

B. Transitional Cases Between Preferences For and Against Particular Types of People and Preferences for Specific Goods and Services: Proxies and Reactions

1. Discrimination Based on Proxy Traits and Stereotypes

Day in and day out, in almost every context, we react to people based on traits they possess which, though immaterial in themselves, we believe to be highly correlated with those traits in which we are primarily interested. We are always using some traits as proxies for other traits. In conclusively presuming for purposes of a particular decision that an individual with a proxy trait possesses the material trait, we stereotype those with the proxy trait.

We could not function without proxies and the stereotypes on which they are based. Proxy traits are not only those traits that correlate with the technical skills required for a particular task, such as having outstanding academic credentials and being qualified for a faculty position at a law school, or, negatively, being female and being unable to play in the National Football League. Proxy traits include all traits that correlate highly or are believed to correlate highly with the traits material to a particular decision. If I conclude that, because a neighborhood is black and poor, I, as a white driver, ought to keep my windows up and doors locked, I am making a proxy decision. If I decide not to hire a young woman because I think it likely that she will quit work soon to have children, or that, given social realities, she is likely to follow her husband if his job is transferred, I am making a proxy decision. If I assume that most of the older people in a Wyoming ranching town restaurant are Republicans and act on that assumption without further investigation, I am making a proxy decision. If I decide not to look for potential employees in locations where I predict the cost of interviewing them exceeds the probability times the extra benefit of finding the most skilled persons at those locations, I am making a proxy decision. If I decide not to date women who are not college-

¹⁸ Physical attractiveness does give its possessors some competitive advantage in the job market, and physically attractive women are especially advantaged relative to homely women. The deformed and grotesque no doubt suffer severe disadvantages. See Note, *Facial Discrimination: Extending Handicap Law to Employment Discrimination on the Basis of Physical Appearance*, 100 HARV. L. REV. 2035, 2036-42 (1987).

educated because it is unlikely that they and I will have much in common, I am making a proxy decision.

The list is endless, the contexts ubiquitous. Indeed, sometimes the trait we think is the material trait for which another trait is a proxy turns out, on reflection, to be a proxy trait for the truly material trait. Thus, certain credentials might be used by an employer as proxies for relevant technical skills, but the skills themselves are only proxies for the ability to increase the employer's wealth, an ability which an employee without the best technical skills might actually possess to a greater degree. Indeed, even the increase in wealth is probably but a proxy for, as well as a means to, an increase in happiness, which the employer might realize to a greater extent were she surrounded by a less productive but more congenial staff.

Proxies correlate with the traits for which they are proxies in three significant ways. First, they may correlate through the laws of nature. Women live longer than men, due in part to biological differences between men and women; being female is thus a proxy for greater longevity for those who set life insurance premiums. Being black correlates highly with certain diseases (e.g., sickle cell anemia) because of biology, and being white correlates highly with others (e.g., skin cancer). Discrimination based upon laws-of-nature proxies may be quite rational because it is cost-justified, even in cases where the correlation between the proxy trait and the material trait is only slightly positive. The cost justification of such proxies is a function of the strength of the correlation and the costs and the improvement of correlation attributable to alternatives to the proxy.

The second significant way in which proxies correlate with their material traits is through the "laws" of human personality and culture: the predicted behavior of persons with the proxy traits. This is the realm not of the laws of physical nature but of psychological and sociological generalizations and predictions. Thus, the employer who fears that a woman is more likely to quit her job than a man because of the desire to have children or the desire to accommodate a spouse's career is drawing upon the sociological generalizations that women in our culture are more likely to engage in childcare functions and more likely to sacrifice their careers for their spouses' than men. These generalizations, though not based on invariant physical laws, may be as or more accurate and confidence-inspiring than physical generalizations. For example, auto insurance companies predict the highway mortality and collision figures for particular groups of drivers with uncanny accuracy

despite the fact that there are no laws of nature compelling speeding, intoxication, or most of the other behavior in question. It is this second category of correlations between proxies and material traits that most people find problematic. This category often evokes the pejorative use of "stereotype" ("Jews are pushy," "blacks are lazy," "women are overly emotional," "whites have no 'soul,'" etc.). This category is frequently labelled "guilt by association" ("Even if most women do quit their jobs to have children, *I* am not going to quit"; "the police should not judge *me* by my long hair, my VW van emblazoned with daisies, and my guitar, even if most people with long hair, VW vans emblazoned with daisies, and guitars are carrying drugs;" etc.).

The third category of correlation between proxies and material traits results from a prediction of how others will react to those with the proxy traits. For example, black police may be no more skilled in policing the ghetto than white police, except that we can predict that the black residents will communicate more openly with the black police and thereby make them more effective than white police. Similarly, male gynecologists may be less effective than female gynecologists of equal technical ability because the patients are more comfortable with the female gynecologists. These "reaction qualifications" are the subject of the next section, so I shall defer any further discussion of them until then.

Proxy discrimination and the stereotypes on which it is based is usually as unproblematic as it is commonplace. Proxy discrimination is quite frequently rational because the underlying stereotypes or generalizations are fairly accurate. Society could not function very well if our generalizations about other people were too frequently inaccurate. Nevertheless, proxy discrimination can be morally troublesome in three ways.

First, irrational proxy discrimination, based upon inaccurate stereotypes or generalizations, is morally troublesome because it imposes unnecessary social costs. So, for the same reason, is proxy discrimination that is irrational, not because it is based on inaccurate stereotypes or generalizations, but because there are more cost/benefit-justified alternative proxies available. Both types of irrational proxy discrimination represent preferences premised on factual errors. And if significant social costs accompany irrational proxy discrimination, it may be morally wrong to engage in it.¹⁹

¹⁹ For various reasons, however, there may be moral rights not to have all morally wrong proxy discrimination legally proscribed. See discussion *infra* Part III.D.1.

Proxy discrimination which is based on accurate generalizations and which is cost/benefit-justified relative to alternative choices exemplifies instrumental rationality. One might think, therefore, that proxy discrimination that is not cost/benefit-justified because it is based on inaccurate stereotypes and generalizations, or because better proxies are available, would not be a major social problem. After all, an instrumentally irrational proxy hurts its user, who then has every reason to switch to more rational proxies once they are discovered. Irrational proxies are a social problem, however, because many of them represent displaced biases. One who realizes that his biases cannot be justified on their own terms, such as one who realizes the invalidity of his judgment that blacks are inherently morally inferior, may, rather than relinquish the judgment fully, merely replace it with a belief that blacks very frequently have trait X, trait X being a perfectly respectable basis for discrimination. Thus, many irrational proxies are the products of bias-driven tastes for certain erroneous beliefs.²⁰ For that reason, the erroneous beliefs underlying irrational proxies, and thus the irrational proxies themselves, being taste- rather than evidence-based, will often prove to be both impervious to contradictory evidence and incorrigible. This is true especially when the costs of error are not very high for the stereotyper. Like the unconscious biases on which they are based, irrational stereotyping and proxy discrimination will be much harder to eliminate through education, at least among those without ample alternative sources of relative self-esteem.

The third troublesome type of proxy discrimination consists of that proxy discrimination that, though based on accurate stereotypes or generalizations, and though cost/benefit-justified, nonetheless has undesirable social consequences. Proxy discrimination based on accurate predictions of the choices of the dispreferred tends to perpetuate the social realities that make the predictions accurate. For example, if women are allowed to drink at an earlier age than men because they are generally more responsible drinkers, men might be reinforced in the attitudes that foster their relative irresponsibility. Similarly, if blacks are dispreferred in employment because the employer has reliable statistics showing that blacks are somewhat more likely than whites to malingering on the job, or to engage in petty theft from the employer, the attitudes that spawn those predicted behaviors might be reinforced.²¹ Even stereotypes

²⁰ For excellent discussions of stereotyping resulting from unconscious biases, see TAJFEL, *supra* note 2, at 143-61; Lawrence, *supra* note 17, at 331-39.

²¹ It is frequently noted that proxy discrimination in employment based on

based on physical laws, such as those relating to life expectancy or susceptibility to various types of illnesses may cause resentment and may reinforce biases and other inaccurate stereotypes. Stereotypes based on psychological and sociological predictions, however, usually prove troublesome even when accurate and otherwise rational. Individuals most likely will resent these stereotypes and view them as insulting and stigmatic ("guilt by association"), especially when these judgments support a conclusion that reflects negatively on moral worth. The same experience may result even when these stereotypes support a morally neutral conclusion which nevertheless disadvantages those whose behavior is predicted.²² Moreover,

accurate and currently cost-justified stereotypes may result in those with the dispreferred proxy traits underinvesting in their human capital if the proxy traits are difficult to change, given that they, as individuals, face collective action problems in changing the general behaviors that make the stereotypes accurate. See Mark Kelman, *Concepts of Discrimination in "General Ability" Job Testing*, 104 HARV. L. REV. 1158, 1160-61, 1232-33 (1991); David A. Strauss, *The Law and Economics of Racial Discrimination in Employment: The Case for Numerical Standards*, 79 GEO. L.J. 1619, 1626-27 (1991). Related to this point is that a proxy discrimination that is based on a currently accurate stereotype that is currently cost-justified may not be dynamically cost-justified: forgoing the use of the proxy, may entail losses in the short term, but may bring about changes in behavior that produce long-term gains that offset the short-term losses. See Kelman, *supra*, at 1160-61; see also Cass R. Sunstein, *Why Markets Don't Stop Discrimination*, 8 SOC. PHIL. & POL'Y 22, 29, 31 (1991). As Sunstein notes:

[I]f there is prejudice and statistical discrimination, and if third parties promote discrimination, there will be decreased investments in human capital. Such decreased investments will be a perfectly reasonable response to the real world. And if there are decreased investments in human capital, then prejudice, statistical discrimination, and third-party effects will also increase. Statistical discrimination will become all the more rational; prejudice will hardly be broken down; consumers and employers will be more likely to be discriminators.

Id. at 31.

Of course, forgoing the use of a good proxy may be dynamically rational socially but not individually, especially where the long term is very long or the short-term losses are quite severe or maldistributed.

²² Stereotypes that support proxy judgments about moral worth—e.g., "people with trait X tend to act immorally or criminally"—not only cause resentment among those people with trait X who do not fit the stereotype, but also tend to produce the very immoral or criminal behavior that justifies the stereotype. See DAVID A.J. RICHARDS, *FOUNDATIONS OF AMERICAN CONSTITUTIONALISM* 279-81 (1989). The problematic status of proxy decision-making based on such stereotypes may help explain the controversy over statistical methods of proof in civil and criminal trials. See, e.g., Judith J. Thomson, *Liability and Individualized Evidence*, LAW & CONTEMP. PROBS., Summer 1986, at 199 (criticizing decisions based on statistical evidence and proposing instead the use of individualized evidence); cf. Richard Schmalbeck, *The Trouble with Statistical Evidence*, LAW & CONTEMP. PROBS., Summer 1986, at 221 (noting a distrust of statistical evidence but recognizing its value if analyzed carefully).

beyond resentment, and beyond the freezing of the underlying social reality mentioned above, such accurate stereotypes may reinforce biases and other inaccurate stereotypes to a much greater extent than stereotypes based on physical laws.

Nonetheless, in assessing whether there is a moral right to make a rational proxy discrimination based on accurate stereotypes, one must always be aware that this discrimination is rational from the discriminator's standpoint. If there are social costs associated with accurate stereotypes in some instances, there are at the same time, individual costs associated with a moral injunction against their use. There is reason to doubt the truth of moral propositions that demand considerable individual sacrifice for the social good.²³ Although the discriminator and those who bear the costs of discrimination are protected by my assumption of a just framework, even within the framework, and surely in setting its boundaries, we should not expect isolated individuals to make considerable sacrifices merely to avoid a greater social cost. Even within an otherwise just framework, social costs should be borne equitably. Of course, the social cost of rational proxy discrimination is simply the aggregation of costs to individuals. Quite frequently, however, the cost to the proxy discriminator of forgoing the use of the proxy will be much greater than the cost any particular individual will bear as a result of proxy discrimination.²⁴

This discussion of the morality of proxy discrimination carries even more force when applied to the legal proscription of proxy discrimination. Given that proxy discrimination reflects a rational attempt to satisfy unbiased and otherwise morally proper preferences, it will be difficult to suppress legally. Moreover, attempts at legal

²³ See KURT BAIER, *THE MORAL POINT OF VIEW: A RATIONAL BASIS OF ETHICS* 203-04 (1958); GILBERT HARMAN, *THE NATURE OF MORALITY: AN INTRODUCTION TO ETHICS* 157-62 (1977); BERNARD WILLIAMS, *ETHICS AND THE LIMITS OF PHILOSOPHY* 77 (1985). But see SHELLY KAGAN, *THE LIMITS OF MORALITY* 386-403 (1989) (arguing that moral demands may be quite severe and discussing the value of being moral and accepting morality's challenges); Peter Railton, *Alienation, Consequentialism, and the Demands of Morality*, 13 PHIL. & PUB. AFF. 134, 160-63 (1984) (suggesting that altering social and political structures can lead to decreased disruption in life presently engendered by quite far-reaching moral demands).

²⁴ Contrast the potential costs to those individuals barred from making a rational proxy judgment that a certain neighborhood is a high-crime one with the potential costs to those persons victimized by that judgment—those with no criminal tendencies whom the discriminators will take special precautions around or avoid entirely. For a hardheaded look at such proxy judgments, see Michael Levin, *Responses to Race Differences in Crime*, 23 J. SOC. PHIL. 5 (1992).

suppression will likely stimulate the invention and use of more ingenious proxies that correlate highly with the forbidden proxies.

2. Reaction Qualifications

Frequently, and particularly when choosing whom to employ for specific jobs, we discriminate on the basis of "reaction qualifications" in addition to technical qualifications. Alan Wertheimer defines reaction qualifications as "those abilities or characteristics which contribute to job effectiveness by causing or serving as the basis of the appropriate reaction in the recipients. Technical qualifications refer to all other qualifications (of an ordinary sort)."²⁵ Wertheimer proceeds to offer some examples of discrimination that are based on reaction qualifications:

- (1) An elementary school principal must choose S or T to teach first grade. S has superior pedagogical skills but has a thick foreign accent which six-year-olds find odd and difficult to understand. The children will actually learn more from T.
- (2) A high school principal must choose U or V to teach a ninth-grade class with serious discipline problems. Whereas U is better trained, he is short and has a high-pitched voice. V is tall, muscular, and has a deep authoritative voice. For this reason, V will have fewer discipline problems and will elicit more learning.
- (3) An advertising agency must choose Y or Z to model swimwear. Although Y's posing technique is superior, Z has the physical attributes (tall, thin, small bustline) which make this swimwear appear more attractive.
- (4) A university health service must choose M (male) or F (female) as a staff gynecologist. M is a superior diagnostician, but many female students feel more at ease with a female gynecologist, will respond better to a female's advice, and will not seek needed medical care from a male physician. There will, therefore, be fewer medical problems if F is chosen.
- (5) A shoe store owner must choose B (black) or W (white) as a salesman. B can fit shoes better, but because many customers are hostile to blacks, W will sell more shoes.
- (6) An appliance store owner must choose H or L as a salesman. H has superior knowledge of appliances but has an aggressive hard-sell personality. L is low-key, and customers in this region (although not in all regions) will buy more from low-key salesmen.²⁶

²⁵ Wertheimer, *supra* note 2, at 100.

²⁶ *Id.* at 100-01. I have omitted Wertheimer's seventh example, choosing a left-handed pitcher to face a predominantly left-handed batting order, because the relevant reaction is arguably not psychological but physical or perceptual.

To Wertheimer's list I would add: (7) B, a black police officer, would like to be assigned to a rural beat, and W, a white police officer, would like to be assigned to the inner city. The department concludes, however, that because of anti-black biases, stereotypes, and aversions among rural whites, W would be more effective than B on the rural beat, and that because of widespread fear of or discomfort around white cops among inner city blacks, B would be more effective than W on the inner city beat. (8) O is a white professional basketball player—the only white on his team—who has been a fan favorite for many years. He is now past his prime and inferior in skill to some black candidates for the squad. Because management believes that the warmth the fans feel toward O plus their possible antipathy toward a team that has no white players add up to higher gate receipts if O is kept on the team, even at the cost of a somewhat worse record, O's contract is renewed. (This last example illustrates not only a reaction qualification, but also the point that skills in business are usually just proxies for making money for the employer.)

After giving his examples of reaction qualifications, Wertheimer continues:

Several observations can now be made. First, it is clear that reaction qualifications are crucial to a wide variety of jobs. Indeed, the entire point of many jobs is to elicit the appropriate reaction. This is true, in part, because many jobs in a modern society involve some form of interpersonal relations—advising, ordering, teaching, selling. . . .

Second, many questions concerning reaction qualifications are essentially empirical: What are the reaction tendencies of the relevant recipients? How do they come about? Can they be changed? Reaction tendencies do clearly vary according to factors such as race, sex, age, class, and region. And whatever their causal basis, at least some reaction tendencies can no doubt be changed by deliberate social action. That said, the reaction tendencies of the relevant recipients are what they are. And they are crucial to job effectiveness. The fact that those tendencies could be different and/or changed should not cause us to deny the present facts or to overestimate their malleability. Some reaction tendencies may be deeply rooted in general developmental patterns or changeable only at great social or individual costs.

Third, to say that recipients' reactions are relevant to job effectiveness is not to say that the employee's actions are irrelevant. Some reaction qualifications do stem from a reaction to a passive characteristic of the employee (e.g., race), but others are elicited

by an employee's actions. It is important to note here that what counts as a technical ability for many jobs is often determined by and cannot be abstracted from the relevant reactions.²⁷

If we take the reactions as given, either because we cannot affect them, at least in the short term, or because we do not believe we should, how should we assess the morality of considering reaction qualifications? Wertheimer argues that reaction qualifications should not be discounted merely because the relevant reactions are based on immutable or passive characteristics, or because the reactions are non-rational.²⁸ On the other hand, reactions that are based on judgments of differential moral worth or inaccurate stereotypes are more appropriately discounted.²⁹ Wertheimer would also be less prone to discount a reaction that is psychologically deep-seated, or that is a personal preference for one's own group's members by a member of a previously victimized group.³⁰ In short, those reactions tend to be most discountable which express intrinsically immoral preferences, such as those reflecting biases and inaccurate stereotypes. Other reactions should not be discounted.³¹

It would be tempting, and analytically tidy, to conclude that the morality of considering reaction qualifications follows straightforwardly from the morality of the reaction preferences themselves. Thus, if a reaction preference were intrinsically immoral because based on a judgment of moral inferiority or on an inaccurate stereotype, considering the reaction as a qualification would itself be immoral. Additionally, if a reaction preference, though not intrinsically immoral, had the kinds of devastating psychological and social effects that would support a conclusion that the preference was immoral, considering the reaction as a qualification would again itself be immoral.

Things are not so tidy, however. First, many otherwise immoral reaction preferences are preferences of individuals who are not fully morally responsible. For example, a bias, aversion, or stereotype held by children that would be immoral if held by an adult may nonetheless substantially impair children's education if not taken into account in hiring teachers. And children cannot be deemed morally responsible for their failure to learn.

²⁷ *Id.* at 101-02.

²⁸ *See id.* at 102-03.

²⁹ *See id.* at 107-08.

³⁰ *See id.*

³¹ *See id.*

Moreover, and more importantly, if we decide we cannot or should not affect the reactions themselves, even though they are immoral, then failure to count the reactions as qualifications will frequently impose costs on parties other than the immoral reactors. For example, failure to assign a black police officer to an inner city beat, where community reaction may make her more effective than a white officer, may result in a higher level of crime, crime that victimizes many persons who do not themselves react to the skin color of police officers.

Thus, even when the reactions are immoral, treating those reactions as qualifications may not be. Indeed, sometimes disregarding immoral reactions may itself be immoral. My tentative conclusion is that the morality of one's treatment of reaction qualifications is not primarily a function of the intrinsic morality or immorality of the reactions. Rather, the morality of the chosen treatment is primarily a function of both considering the gravity and the distribution of the social effects of acknowledging reaction qualifications and the gravity and distribution of the social effects of not acknowledging them.³²

*C. Preferences for Goods, Services, and Personal Relations
That Entail Preferences for Particular Kinds of People*

I began my survey of discriminatory preferences with primary preferences for and against certain people. I then moved to secondary preferences for and against certain people based on proxy or reaction qualification considerations. These preferences, though not primary, were intermediate between preferences for goods and services and primary preferences for and against certain kinds of people. The proxy relationship or the reaction qualification frequently reflected or reinforced primary preferences for and against certain kinds of people. I now move to preferences purely for goods, services, and personal relations that nonetheless entail preferences for and against certain kinds of people.

By far the most common type of discriminatory preference is the preference for people with traits that are technical qualifications for

³² To the limited extent that the morality of counting reaction qualifications depends upon the intrinsic morality/immorality of the reactions, it is because it is less problematic to expect the reactors to give up intrinsically immoral preferences than it is to expect them to give up intrinsically benign ones. I owe this point to Ken Greenawalt. Letter from Ken Greenawalt to Larry Alexander (Nov. 1991) (on file with the author).

specific tasks and relationships. This type of preference views the traits in question not as primarily valuable in themselves, but valuable as means to production, entertainment, and other ends. If we enjoy good basketball, then we value—in that context—those people with the traits conducive to playing good basketball. If we value skilled neurosurgery, then we value—again, in that context—the people with the traits conducive to skilled neurosurgery. If we value good but inexpensive widgets, then we value those who are most efficient at producing them. If we value humor in a companion, then we seek as companions those who possess a good sense of humor.

In any large, technologically advanced society that has a sizable realm of individual freedom, including market freedom, there exists a numerous and diverse set of esteemed traits and skills. Many who rank high on one scale of esteem will rank low on others, and vice versa. Skilled neurosurgeons may make lots of money but may not be valued highly as companions because of dour personalities. Unskilled workers may be great companions. Nonetheless, it would be naive to assume that equally esteemed traits and skills are distributed in such a way that everyone scores as well as everyone else when the scores on the scales measuring these traits and skills are summed. And of course, with respect to traits and skills that are the subjects of economic reward, a free market clearly does not distribute monetary rewards equally.

Should we be troubled by discrimination and inequality that follows from free choices in the economy and elsewhere? Recall that I am prescinding the moral analysis of discrimination from other matters of social justice such as whether there is a just distribution of wealth. Recall also that I am assuming that having a large realm of freedom, including a substantially free market, is not per se inconsistent with justice. Is there something about certain distributions of differential esteem and inequalities of material reward that is morally problematic when similar distributions along different axes would not be morally problematic?

For purposes of analysis, I am going to break the category of preferences for goods, services, and relationships into two sub-categories: preferences that are “tainted” by associations with biases, stereotypes, and dubious ideologies; and preferences that are completely untainted by such associations. The former is more closely connected with the preferences I have already considered, so I shall address it first. The latter is at the opposite end of the

preferences spectrum from the biases with which I began, so I shall address it last.

1. Preferences for Goods, Services, and Personal Relations Tainted by Associations with Biases, Stereotypes, and Ideologies

a. *Generally*

Many of our preferences for goods, services, occupations, and the like can be traced historically to biases, stereotypes, ideologies, and other dubious tribalisms and to the social structures to which they gave rise. For example, our taste for certain types of music such as "the blues" may be traceable to the structures and attitudes of the Jim Crow era. Similarly, our standards for what counts as female beauty for purposes of beauty pageants and selecting models and starlets may reflect a preference for Caucasian characteristics traceable to racial animosities and pride. Our preference for distinctive men's and women's clothing fashions rather than unisex fashions, a preference that in turn makes women more technically (as well as reactively) qualified to model women's fashions and men more technically (and reactively) qualified to model men's fashions, may be rooted in gender ideology. Male aversions to becoming secretaries ("women's work") and female aversions to certain "male" occupations are commonplace and obvious examples of this type of preference.³³ Many other preferences which seem on their face to be unrelated to these tribalisms, however, may in fact be related. We may prefer steeply pitched roofs in houses because of long forgotten associations with particular religious sects and their churches. We may prefer certain sports because of connections with past ideologies now buried in the collective subconscious.

These "tainted" preferences will vary considerably in terms of how strong, how deeply rooted, and how autonomous from their tainted origins they now are. Moreover, the discrimination now brought about by such preferences will not necessarily parallel past discrimination or reflect the normative hierarchy that gave rise to it. Although black models may be relatively disadvantaged by "white" standards of beauty, the male aversion to "women's work" tends to raise the relative wages of secretaries (predominantly

³³ There is a burgeoning body of feminist literature focusing on "genderized preferences." See, e.g., Christine A. Littleton, *Reconstructing Sexual Equality*, 75 CAL. L. REV. 1279, 1284-85 (1987).

females) and depress the relative wages of construction workers (predominantly males). Although black musicians may find it much harder psychologically than white musicians to fulfill requests for "Dixie," they may find it relatively easier to satisfy white audiences' tastes for jazz and rhythm and blues. The point here is not empirical, but conceptual: the impact of a tainted preference is logically independent of the impact of those past preferences that make it "tainted."

Should anyone feel aggrieved by being disfavored for a job or some other benefit because of a preference or taste whose historical origin is some immoral bias, ideology, or social condition? If we currently value jazz and those who play it well, does it matter that, but for past racism, we would not have such a taste? Should we feel guilty that our standards of human beauty are almost certainly tribal and not universal? Should our attitude towards secretaries' salaries be affected by the knowledge that, but for a now widely discredited ideology of gender roles, secretaries' salaries might well be lower (or higher, if male aversions are more than offset by female preferences)? Although some of these "tainted" preferences might be quite discrete, quite shallow psychologically, and quite malleable, most will be difficult to identify, to disconnect from other preferences, and to excise. Moreover, because the tainted preferences do not necessarily produce current discrimination that tracks the immoral discrimination that provides the "taint," it is difficult to see how these preferences and the discrimination to which they give rise are morally distinct from untainted preferences and the discrimination to which they give rise.

b. *The Special Case of Unconscious Bias*

Before turning to ordinary, untainted preferences and the discrimination they produce, I want to discuss separately one special category of tainted preferences: preferences reflecting currently operative but unconscious biases. The category has come under discussion recently in connection with the Supreme Court's requirement that racially discriminatory intent be proved to establish a violation of the equal protection clause.³⁴ My concern with

³⁴ See *Washington v. Davis*, 426 U.S. 229, 242 (1976). For a discussion of the relation of unconscious bias to the Court's discriminatory intent standard, see Lawrence, *supra* note 17; David A. Strauss, *Discriminatory Intent and the Taming of Brown*, 56 U. CHI. L. REV. 935, 960-62 (1989).

unconscious biases is, of course, different. I want to ask if such biases, or rather the tainted preferences by which they are manifested, have the same moral status as preferences reflecting conscious biases, or whether instead the unconsciousness of the biases makes such preferences morally no different from other tainted preferences that have become autonomous from their tainted origins.

The type of preference I am concerned with is one that disproportionately negatively affects members of a particular group but, at a conscious level, is neutral with respect to that group. Nonetheless, the preference is anything but neutral, for it will be abandoned in favor of other "neutral" preferences that negatively affect the group if and when it ceases to hurt and begins to benefit the group in question. The explanation for this dynamic is that the discriminator is biased against the group, though the bias is unconscious.³⁵

For example, someone might be an avid Boston Celtics fan when the team is predominantly white. When blacks come to dominate the team in terms of numbers, the fan finds that he just cannot get excited about the Celtics any more, and he switches his interest to the all-white Boston Bruins hockey team. He does not attribute this change in preference to racial bias, for he does not believe he is biased and rejects all biased judgments at the conscious level. He is, however, unconsciously biased, which means that if the Celtics became predominantly white again, or if the Bruins became predominantly black, his allegiances would shift back to the Celtics.³⁶

As thus described, unconscious biases are best thought of as dispositions,³⁷ unaccompanied by conscious biased judgments,³⁸ to prefer and disprefer particular goods, services, and people when—and because—such preferences have certain disproportionate group impacts. What shall we say about their moral status?

³⁵ What I have to say about unconscious biases fully applies to unconscious stereotypes that, were they conscious, would produce irrational discrimination because of their degree of inaccuracy and the availability of superior proxies. *See supra* text accompanying note 19.

³⁶ For other examples of preferences reflecting unconscious biases, see Lawrence, *supra* note 17, at 348-49.

³⁷ See MICHAEL MOORE, *LAW AND PSYCHIATRY: RETHINKING THE RELATIONSHIP* 322-42 (1984).

³⁸ I realize that much, much more can be said about the phenomenology of bias and stereotype-driven discrimination and the roles that judgments regarding moral worth and likely characteristics play in discriminatory actions, a realization based largely upon conversations with Ken Simons. Interviews with Ken Simons, Professor of Law, Boston University School of Law, in Boston, Mass. (Nov. 1991).

The case for treating unconscious biases as morally equivalent to conscious biases and unlike other tainted preferences is surely a weighty one. The preferences of the unconsciously biased are not simply tainted because of their origins. Unlike tainted preferences which have become autonomous from their tainted origins, the preferences of the unconsciously biased track the preferences of those who are consciously biased. Dispreferred individuals will likely find them just as stigmatic and insulting as consciously biased preferences.³⁹ Moreover, such preferences should be relatively shallow psychologically, because they rest on biases that are rejected at the conscious level. Once made aware of their unconscious biases, the discriminators are quite likely to abandon the tainted preferences.

Despite the strength of the case for treating unconscious biases as on a moral par with conscious biases, I am unable to reach any firm conclusion on the matter, for there is also a case for treating unconscious biases no differently from other tainted preferences and thus—if one accepts my conclusion about tainted preferences—no differently from ordinary preferences that have disproportionate group impact. First, it is not clear that a preference, as opposed to a judgment, can be morally defamatory. That I prefer basketball to football does not imply that I regard those who play football as inferior human beings. Preferences do imply value judgments—that what is preferred is of greater value to the preferer than alternatives—but not all such value judgments are capable of universalization. That I regard something as valuable does not mean that I regard it as valuable to you, even if I accept that value judgments are objective and that values are grounded in reality. The preferences of the unconsciously biased, *qua* preferences, do not automatically entail defamatory judgments about moral worth.

More significantly, it is not at all certain that one who becomes aware that a preference structure reflects unconscious biases he consciously disavows has any reason to and thus can be predicted to give up the preference structure. The Bruins fan, when made aware of the unconscious racism that has caused him to prefer hockey to basketball, may still find that he prefers hockey to basketball. There has been a great deal written about the morality of satisfying preferences that are “inauthentic,” adaptive to an immoral reality, or the products of conditioning to accept an incorrect ideology.⁴⁰

³⁹ See Lawrence, *supra* note 17, at 352-54.

⁴⁰ See RICHARD B. BRANDT, A THEORY OF THE GOOD AND THE RIGHT 88-129 (1979);

These discussions tend to assume—incorrectly, I believe—that “cognitive psychotherapy” will cause us to reject those preferences of ours that are traceable to such morally problematic causes.⁴¹

This premise may be true where the preferences rest entirely on incorrect or questionable beliefs that are now exposed as such, as when a special preference or aversion rests upon an incorrect belief implanted in childhood. Nevertheless, it is more dubious in the cases considered here, where the bias against blacks is consistent with preferring many things besides hockey. In other words, the bias explains why hockey was eligible for preference, relative to basketball, but it does not explain why hockey is preferred over other sports dominated by whites.

My tentative conclusion is that when unconscious biases result in personal aversions—where what is dispreferred is no more than association with the group targeted by the unconsciously biased discrimination—the aversion has the same moral status as a conscious bias.⁴² It will be experienced by the dispreferred as they would experience conscious bias. Furthermore, it should be easy to extirpate once its source in unconscious bias is exposed, because the bias itself is consciously rejected, and because the bias fully determines the aversion.

On the other hand, when the preference stemming from unconscious bias is for some good or service and is not merely an aversion to certain people, my weak intuition is to treat it like other tainted preferences and thus like ordinary preferences. The only difference between preferences reflecting unconscious biases and other tainted preferences is that the tainted attitudes are still active in the former case. Although active, they are nevertheless unconscious, and that is significant: unconscious “judgments” are judgments in only an

JON ELSTER, *SOUR GRAPES* 43-66 (1983); DIANA T. MEYERS, *SELF, SOCIETY, AND PERSONAL CHOICE* 59-75 (1989); Cass R. Sunstein, *Legal Interference with Private Preferences*, 53 U. CHI. L. REV. 1129, 1152-58 (1986).

⁴¹ See BRANDT, *supra* note 40, at 11. There is a growing body of philosophical literature on second-order preferences, or preferences about preferences, the progenitor of which is Harry G. Frankfurt's *Freedom of the Will and the Concept of a Person*, 68 J. PHIL. 5 (1971). See also Paul Benson, *Autonomy and Oppressive Socialization*, 17 SOC. THEORY & PRAC. 385, 391-94 (1991) (considering Frankfurt's analysis of higher order desires with respect to the oppressive socialization of women); John Christman, *Autonomy: A Defense of the Split-Level Self*, 25 S. J. PHIL. 281, 283-90 (1987) (analyzing higher and lower order desire theory of autonomy, including criticism and defenses); Thomas E. Moody, *Liberal Conceptions of the Self and Autonomy*, in *FREEDOM, EQUALITY, AND SOCIAL CHANGE* 94, 94-98 (Creighton Peden & James P. Sterba eds., 1989) (analyzing Frankfurt's hierarchical theory of autonomy).

⁴² See *supra* text accompanying note 17.

extended metaphorical sense. It is also not clear that one can be morally defamed by such a behavioral disposition any more than one can be defamed by a robot programmed to act as if it loathes you. In addition, unconscious biases underdetermine the preferences that manifest them. While we can say that the bias itself is wrong, we cannot say the same of the preferences. We can attempt to make unconscious biases conscious with the hope of eradicating them; but the preferences might very well persist long after the bias disappears.

2. Preferences for Goods, Services, and Relations That Differentially Affect Various Groups

Our preferences for particular goods, services, and relations and the finitude of our resources mean that, given a realm of freedom in a society of sufficient size and technological attainment, some people will command higher salaries, will be more esteemed, and will have more satisfying lives than others. Our preference for NBA basketball over archery means that those with the skills required for NBA basketball will be financially better off than skilled archers. I may be the only person in the United States capable of making a good sixteenth century lute; but if no one wants a good sixteenth century lute, and if no one esteems the skills that go into making one, I will neither be rewarded nor admired for my rare talent.

The inequalities of wealth and esteem produced by a realm of freedom to express preferences lead to, of course, inequalities among John, Jane, Joe, and Joan. They can also be viewed, however, as inequalities among groups. We are each members of an indefinite number of groups. For example, I am a member of the following groups: the American-born; males; Jews; persons under five feet nine inches; persons over five feet seven inches; Californians; Padres fans; lawyers; law professors; ex-Texans; persons over 135 pounds; the blue-eyed; married persons; parents; fathers of one girl and two boys; graduates of Eastern colleges; and so on, and so on. No matter along what axis we group human beings, some of those groups will fare better than others in terms of income, esteem, and satisfaction. For example, the group "males" *qua* group has a higher per capita income than the group "females."⁴³ Within the group "women", white women generally enjoy higher incomes than black women.⁴⁴

⁴³ See BUREAU OF THE CENSUS, U.S. DEP'T OF COMMERCE, STATISTICAL ABSTRACT OF THE UNITED STATES 457 (1991) (table 736 showing that median money income of year-round full-time workers in 1989 was \$28,605 for men and \$19,643 for women).

⁴⁴ See *id.* (showing that the median money income of year-round full-time female

On the other hand, the group "ex-Texans", or "people under five feet nine inches," may not be doing particularly well in relative terms. Surely, Padres fans are currently not faring well.⁴⁵

When it arises, not from group biases, ideals, ideologies, aversions, reactions, or stereotypes, but solely as a product of ordinary preferences for goods and services and the distribution of skills, physical attributes, attitudes, ambitions, and so forth, does inequality among any of these groups have any moral significance? It is tempting, and perhaps correct, to argue that these group inequalities—disparate group impacts—are of no moral import. Any set of preferences produces inequalities between the relatively preferred and the relatively dispreferred. Therefore, why should we attribute any moral significance to some groups whose membership correlates positively with the group consisting of all relatively dispreferred individuals, when we don't attribute moral significance to the group of relatively dispreferred individuals itself? Put differently, if there are ten individuals with different incomes, why should we feel more morally troubled if the bottom five are predominantly black and female and the top five are predominantly white and male than if blacks and females are distributed more evenly, especially if there is no present bias, inaccurate stereotyping, and so forth? We have ten individuals with different incomes in either case. Furthermore, are we not inappropriately reifying the groups when we assume that they are affected by these inequalities in some way apart from how individuals are affected *qua* individuals?⁴⁶

workers in 1989 was \$19,873 for white women and \$17,908 for black women).

⁴⁵ The Padres finished the 1992 baseball season one game over .500 and barely in third place in the National League Western Division. See N.Y. TIMES, Oct. 6, 1992, at B10 (table showing final National League standings).

⁴⁶ See Michael E. Levin, *Is Racial Discrimination Special?*, 15 J. VALUE INQUIRY 225, 227-28 (1981); see also Walter E. Williams, *False Civil Rights Vision and Contempt for Rule of Law*, 79 GEO. L.J. 1777, 1782 (1991). For examples of such group reification, see Maxine N. Eichner, Note, *Getting Women Work That Isn't Women's Work: Challenging Gender Biases in the Workplace Under Title VII*, 97 YALE L.J. 1397, 1401-14 (1988); Martha T. McCluskey, Note, *Rethinking Equality and Difference: Disability Discrimination in Public Transportation*, 97 YALE L.J. 863, 878-80 (1988). Both authors treat economic equality between men and women *qua* groups, but not equality among individuals or among groups such as the skilled and unskilled, as a moral desideratum that can trump efficiency, employer interests, etc. See also T. Alexander Aleinikoff, *A Case for Race-Consciousness*, 91 COLUM. L. REV. 1060, 1062-81, 1091-95 (1991); Mary E. Becker, *Prince Charming: Abstract Equality*, 1987 SUP. CT. REV. 201, 206-12; Sunstein, *supra* note 21, at 33-36; Cass R. Sunstein, *Three Civil Rights Fallacies*, 79 CAL. L. REV. 751, 771-73 (1991) [hereinafter Sunstein, *Three Civil Rights Fallacies*].

Two different reasons are given for the moral significance of the disparate group impact of ordinary preferences. One points to the present social consequences of disparate impact, and the other to both those present social consequences and to past wrongs.

a. *Disparate Impact and Present Social Consequences*

Although we are each members of an indefinite number of groups corresponding to the various characteristics we share with some but not all other persons, it is a psychological and sociological truth that membership in a few of those groups is more central to our sense of who we are than membership in most of them. If I personally am dispreferred for a particular job or as someone's companion, I will feel the sting of personal rejection as well as the loss of the opportunity denied. My personal preferences to have the job, to be the companion, or to have esteemed traits will be unsatisfied. Beyond my personal defeat, however, I will generally be unconcerned with who in fact gets the positions and esteem I sought. For example, I will be unconcerned with whether the job I sought unsuccessfully was awarded to a white or to a black, to a male or to a female, to a Jew or to a Protestant, to a five foot eight inch Padres fan or to a six foot three inch Dodgers fan.

On the other hand, there are people who are concerned with how others within certain groups are faring. Many blacks care about how other blacks are doing in the job market. If blacks as a group are doing poorly relative to other racial groups, they feel bad regardless of how they as individuals are doing.⁴⁷ The same attitude is true of many women towards women as a group, many ethnic group members toward their ethnic group, and many religious group members toward their religious group. And it is almost universally true of the attitude we take toward our family and group of close friends. Moreover, in a variety of contexts we have preferences regarding group success with respect to groups of which we are not members. For example, many whites take pleasure in the success of blacks as a group; many men take pleasure in the success of women as a group; and avid Padres fans live and die with the fate of the team, and only derivatively with the fate of the fans.

Close identification with particular groups not only produces these vicarious pleasures and pains, satisfactions and frustrations, but also produces subtle and not so subtle effects on aspirations and

⁴⁷ See Kelman, *supra* note 21, at 1240-43; Strauss, *supra* note 21, at 1629-30.

motivations. Blacks' identification with other blacks, coupled with the disproportionately high number of successful black athletes and the disproportionately low number of successful black academics, may result in a disproportionate number of young blacks striving to become athletes rather than academics. Similarly, women may opt for childbearing over corporate management because they identify with other women and see so few in the latter role and so many in the former.⁴⁸ Finally, group consciousness produces political demands on behalf of the groups that exhibit this phenomenon. These demands are in turn a source of social strife and costs regardless of their moral legitimacy.

There are a number of questions regarding the moral significance of these costs occasioned by disparate group impact. First, are they different in kind or in legitimacy from extreme individual reactions to being dispreferred? ("If you don't hire me for this job, I'll kill myself, I'll incite a riot, I'll be psychologically crippled, etc.") Second, should failure to satisfy the external preference for how others are treated⁴⁹ be viewed as a cost of satisfying ordinary preferences for purposes of morally assessing the satisfaction of those ordinary preferences? Third, are these social costs of the disparate group impact of ordinary preferences properly regarded as costs of the preferences, or are they properly regarded as costs of the underlying biases, aversions, commitments, and stereotypes that lead to the group identifications?⁵⁰

There is, finally, a social cost of disparate impact that stems neither from reifying groups nor from personal or vicarious identification with groups, though it is connected with similar phenomena. Adverse disparate impact associated with morally neutral preferences will in some cases reinforce or produce anti-group biases, aversions, and inaccurate stereotypes held by others. If I prefer to hire the ablest legal theoreticians for my law faculty, and very few of those available happen to be black, one effect beyond the effects on individual disappointed black applicants, and beyond the effects on other blacks (and non-blacks) who identify with the fortunes of blacks as a group, are the effects on (mostly) non-blacks in reinforcing or spawning, for example, beliefs in black

⁴⁸ See Sunstein, *Three Civil Rights Fallacies*, *supra* note 46, at 771.

⁴⁹ See DWORKIN, *supra* note 8, at 231-39.

⁵⁰ These social costs are, in fact, clearly joint costs; but morally they are perhaps assignable only to the group identification side of the ledger rather than to the preferences for goods and services side. Alternatively, should they be morally assigned to the cheapest cost avoiders? If so, who are the cheapest cost avoiders?

moral or intellectual inferiority. In assessing the morality of hiring the ablest legal theoreticians, must I consider these latter effects in addition to (or instead of) the others?

b. *Disparate Impact, Present Social Consequences, and Past Wrongs*

Suppose we add another consideration to the previous discussion of disparate group impact and its present social consequences: the reason many ordinary preferences result in an adverse disparate impact on certain groups is because these groups, or at least many of these groups' members, have been the victims of past wrongs. Should this in any way affect the moral assessment of these ordinary preferences or, more precisely, acting on these preferences?

For example, black Americans and American women surely suffered numerous wrongs in the past. Those wrongs have produced, among other things, effects on the distribution of job skills, aspirations, and motivations among present-day blacks and women. Without these past wrongs, it is highly likely that the job skills, aspirations, and motivations of blacks and women would be more like the job skills, aspirations, and motivations of present day white males. Predictably, given the present distribution of skills, etc.—among blacks and whites, females and males—and the present hierarchy of preferences for goods and services, whites and males achieve disproportionately higher income and status than blacks and females.

In assessing whether the genesis of present-day skills, etc., in past wrongs changes the moral analysis of preferences that produce disparate group impacts, the reader should again bear in mind that I am bracketing all issues of distributive justice and just institutions, political rights, etc., beyond discrimination. Included in these bracketed issues of background justice are those relating to whether, and to what extent, reparations are due for the past wrongs as a matter of justice. The case for reparations beyond wrongs narrowly circumscribed both in time frame and in type is quite problematic.⁵¹ In any event, I want to inquire whether, beyond reparations,

⁵¹ See, e.g., Levin, *supra* note 46, at 231 (noting the virtually endless number of wrongs that might be eligible for reparations); Glenn C. Loury, *Affirmative Action: Is It Just? Does It Work?*, in THE CONSTITUTIONAL BASES OF POLITICAL AND SOCIAL CHANGE IN THE UNITED STATES 109, 124-31 (Shlomo Slonim ed., 1990) (pointing out the difficulties in identifying proper beneficiaries and payors of reparations); Christopher W. Morris, *Existential Limits to the Rectification of Past Wrongs*, 21 AM. PHIL. Q. 175, 175-78 (1984) (highlighting personal identity problems in imagining how

past wrongs have any bearing on the morality of present disparate group impact which results from acting on ordinary preferences. It is best then to assume, for purposes of this inquiry, that all claims for reparations have been honored to the extent required by justice.

It is difficult to see how past wrongs affect the assessment of disparate group impact. Many individuals have been affected in terms of their present skills, etc., by past wrongs they or their ancestors suffered, even though they are not members of any group that is conscious of itself as a victimized group aggrieved by disparate impact. (Many unrepaired wrongs have victimized and presently affect white male Protestants. Indeed, many wrongs of the type I am concerned with in this article—discrimination based on bias, etc.—have caused the victimized group to redouble its efforts to succeed and have caused it to do so disproportionately well given present preferences for goods and services. Asians and Jews may be good examples of groups that have disproportionately succeeded not only despite, but perhaps because of, past discrimination. Particular individuals in those groups, however, may suffer present competitive disadvantages because of past wrongs.) Such individuals are, of course, frequently the proper recipients of reparations from wrongdoers, but generally we do not believe that it is wrong to disprefer, say, in the context of picking a surgeon, one whose hands were villainously cut off as she was studying to be a surgeon, and who would have been the best surgeon but for the wrong she suffered.

This seems to suggest that our concern for past wrongdoing is exhausted by the subject of reparations. It does not appear to bear on the morality of acting on present ordinary preferences.

c. The Special(?) Case of Preferences for Traditional Ways of Doing Things

As a final case, assume a preference that has a disparate impact on socially significant groups that have suffered past wrongs, and that is also tainted and perhaps based in part on others' reactions. The best examples—and surely ones that are currently engendering considerable controversy—are preferences for certain traditions. In hiring the president of an Ivy League college, the trustees may prefer someone who fits the traditional image of the school, someone, say, with the looks and the enunciation of a John Houseman or John

present day persons would have fared but for past discrimination).

Gielgud. This preference may well be tainted as I have defined tainted. It surely will disprefer blacks, women, and recent immigrants (except those from the British Isles). And it may well be based on nothing more "objective" than the faculties', students', and alumni's love of the tradition it taps into for its own sake. The same may be true of myriad other traditions—from traditions in the workplace that, say, many women regard as uncomfortable or harassing, to regional or local traditions that certain groups don't fit into well for emotional, physical, or other reasons, to even purely academic traditions regarding scholarly and teaching styles and curricular preferences that might disproportionately negatively impact certain groups.⁵²

The question raised by preferences for traditional ways is whether combining taint, disparate impact, past wrongs, and perhaps reactions produces preferences that are inherently illegitimate even if none of those attributes singly does so. I cannot see how. I have just argued that disparate impact, even when it affects groups that have suffered past wrongs and affects them because of those wrongs, does not in itself make preferences illegitimate.⁵³ The reactions at issue in tainted preferences for traditions are themselves simply tainted preferences for traditions. Therefore, the question is whether tainted origins produce illegitimacy in preferences that result in disparate impact. Since I have also argued that tainted preferences are, in most cases,⁵⁴ morally equivalent to ordinary preferences, the answer must be "no."

The legitimacy of preferences for traditional ways is not just a matter of repudiating various possible sources of illegitimacy. There is also a positive case. Our traditions and our preferences for them in large part define who we are both individually and as a community. All traditions contain some tainted history and disparately impact some groups. Thus, to ask people to repudiate such preferences is to ask them to create their preferences and thus themselves *ex nihilo*.

⁵² The attacks on academic traditions will be quite familiar to most readers of this article. See, e.g., Richard Delgado, *Approach-Avoidance in Law School Hiring: Is The Law a WASP?*, 34 ST. LOUIS U. L.J. 631, 634-41 (1990). *Academic Questions*, the journal of the National Association of Scholars, regularly chronicles the assaults on academic traditions based on charges of taint and disparate impact. Nonacademics will probably be more familiar with attacks on school nicknames—for example, the Stanford Indians (now the Cardinals) or various Southern high school teams called the Rebels.

⁵³ See *supra* text accompanying notes 49-51.

⁵⁴ The exceptions are cases of aversions to people based on unconscious biases. See *supra* text accompanying note 42.

Ironically, many of those who inveigh against preferences for traditional ways also attack liberalism for creating alienated, deracinated individuals, deprived of needed community and tradition.

III. WHEN IS ACTING ON CERTAIN PREFERENCES MORALLY WRONG AND WHAT SHOULD BE DONE ABOUT IT?

The taxonomy of various discriminatory preferences is now complete. In this Part, I wish to offer some very tentative conclusions about the moral status of those preferences and about the implications of my moral assessment for legislation concerning private discrimination. I have raised most of the considerations that bear on moral status in the process of discussing each type of discriminatory preference, and I have at least hinted at what moral conclusions I think should be drawn. Now I want to synthesize these separate and incomplete moral evaluations, even if the general moral picture I paint is very fuzzy and offered with less than total conviction.

First, I should remind the reader once again where the moral inquiry is situated with respect to other moral considerations raised by private discriminatory acts. I am assuming for the limited purposes of this inquiry that the society within which the discriminatory acts occur is basically just, and that no one is, by virtue of these discriminatory acts, deprived of those basic resources and rights to which all persons are entitled as a matter of justice. I am thus prescinding the analysis of discrimination from the more general problem of social justice. I am further assuming that even where a person has a moral right to discriminate, such discrimination can be morally criticizable (i.e., that one can have a "right" to do "wrong"). Moreover, and perhaps most crucially for legal audiences, I am assuming that if discrimination is morally wrong, then, though one may have a moral right to engage in it, it is possible that one may not. If one lacks a moral right to engage in morally wrong discrimination, the case for legal prohibition is open, though not by any means conclusive.

Thus, my inquiry rests on assumptions that are inconsistent with the Right, which treats all preferences as exogenous givens and concerns itself with establishing the moral boundaries/entitlements within which there is a moral right to act on any and all preferences, from the most elevated to the most hateful. My inquiry also rests on assumptions inconsistent with the Left, which, like the Right,

treats all preferences as exogenous givens, though in the Left's case it is because only the distributional pattern of results in terms of wealth, "meaningful" work, etc., matter, not the preferences that produce the distributional pattern. Finally, I am assuming many preferences are revisable on the basis of moral reflection and are not all incorrigible facts, like hunger, with which persons must contend but cannot affect.

Second, it should be clear from the discussion in Part II that notions like "equality" or "equality of opportunity" will not be helpful in assessing discriminatory preferences, at least given the assumptions that frame my inquiry. I share the view expressed by others that "equality" itself is an empty concept for purposes of normative criticism.⁵⁵ Even if that view is incorrect as a general matter, it is surely correct when distinguishing among forms of discrimination on moral grounds, since all discrimination produces inequality between those preferred and those dispreferred. Similarly, the concept of "equality of opportunity" will not aid in assessing the morality of discrimination, because the concept can mean anything from formal legal equality ("it is legally permissible to hire blacks as well as whites") to equality of result ("everyone must get the same jobs, the same working conditions, and the same income; thus, private discrimination of any type is inconsistent with equality of opportunity").⁵⁶ Any decision reached regarding which discrimination is morally permissible and which is impermissible produces a certain conception of equality of opportunity, but it is not—because logically it cannot be—the product of that empty concept.

I think that the most perspicacious way of analyzing the morality of discriminatory preferences is first to divide those preferences into the intrinsically immoral and intrinsically benign. Intrinsically immoral preferences are those that are not held by fully informed and rational persons and that produce unnecessary harms to others. Intrinsically benign preferences may be immoral, but if they are, it is a result of extrinsic, contingent conditions.

⁵⁵ See Larry Alexander, *Constrained by Precedent*, 63 S. CAL. L. REV. 3, 9-13 (1989); Larry Alexander, *Striking Back at the Empire: A Brief Survey of Problems in Dworkin's Theory of Law*, 6 L. & PHIL. 419, 426-31 (1987); Peter Westen, *The Empty Idea of Equality*, 95 HARV. L. REV. 537, 578 (1982).

⁵⁶ See Christopher Jencks, *What Must Be Equal for Opportunity to Be Equal?*, in EQUAL OPPORTUNITY 47, 71-73 (Norman E. Bowie ed., 1988); Westen, *supra* note 55.

A. *Intrinsically Immoral Preferences*

Biases—except for those reflecting close personal ties that are so central to one's identity they amount to "biases" in favor of one-self⁵⁷—are paradigmatically intrinsically immoral. Biases rest on erroneous judgments of others' inferior moral worth, and they insult and produce justifiable resentment in the dispreferred, harms that are unnecessary additions to the inevitable harms of being dispreferred. Preferences based on ideologies regarding role ideals that cannot be defended and that are experienced as oppressive by the dispreferred also fall into the camp of the intrinsically immoral.

On the other hand, ordinary preferences for goods and services, regardless of their impact in terms of income and status on various groupings of individuals, and regardless of their impact on persons who have suffered past wrongs, are intrinsically benign.⁵⁸ Such preferences express no defamatory judgment about the dispreferred, only the judgment that the discriminator prefers some trait the dispreferred lack over traits they possess. One cannot assume both that a realm of free choice is morally permissible and, at the same time, that ordinary preferences and the judgments they express are impermissible.

Although this will no doubt be more controversial, I also contend that tainted preferences are intrinsically benign. Perhaps my view is affected by my difficulties in imagining a set of preferences in which all historical tribalisms that have made particular cultures what they are never existed. Our tastes in art, music, sports, etc., are so heavily culturally influenced, and the culture itself is so much a product of morally unrespectable views now discarded, that the counterfactual world of untainted preferences lies beyond my limited powers of imaginative reconstruction. In any event, a tainted preference can become autonomous from its disreputable origin, and when it does so, I believe it will withstand the test of cognitive psychotherapy—endorsement by the clearheaded and fully self-aware agent⁵⁹—that biases and inaccurate stereotypes will fail. More controversially, I also believe that tainted preferences must be

⁵⁷ See *supra* text accompanying notes 10-11.

⁵⁸ By calling a preference "intrinsically benign," I mean only that it is not intrinsically immoral; intrinsically benign preferences include both the morally neutral and the morally positive, and in either case particular examples of such preferences may be ultimately adjudged immoral because of extrinsic considerations.

⁵⁹ See BRANDT, *supra* note 40, at 11.

deemed intrinsically benign even when they are not autonomous from their biased origins.⁶⁰

Proxy discrimination, when use of the proxy trait is a cost/benefit justified means of satisfying ordinary preferences, is also intrinsically benign. This is so regardless of whether the proxy trait correlates with the trait for which it is a proxy as a matter of physical laws or sociological and psychological generalizations.

Proxy discrimination based upon inaccurate and usually bias-driven stereotyping are intrinsically immoral for the same reasons as are the biases with which they are intimately linked. This is surely true once the inaccuracy of the stereotype or lack of cost-justification is revealed to the proxy discriminator. On the other hand, proxy discrimination that is unknowingly premised on inaccurate stereotypes or is otherwise not cost-justified, but is not bias driven, is best regarded as intrinsically benign.

Consideration of reactions as part of assessing qualifications is also intrinsically benign, despite the fact that it might often be immoral for various other reasons. From the discriminator's standpoint, consideration of reactions is, like proxy discrimination, merely part of finding the best person for the job, which is itself just an ordinary preference. The discriminator is not expressing a biased or distorted judgment about the dispreferred, but is merely rationally seeking the least costly means of satisfying what are often perfectly respectable preferences. (Consider as an illustration the preferential hiring of black police officers to satisfy the preference for improved police protection in areas where reactions to white officers will hinder their performance.) Although the reactions may often express intrinsically immoral preferences, taking them into account does not.

The final, and for me the most difficult to assess, discriminatory preference is the personal aversion to people with particular traits, such as membership in a certain racial or ethnic group or possession of some physical characteristic, inborn or acquired. On the one hand, these personal aversions are closely connected with biases and inaccurate stereotypes and commonly arise from them. Moreover, they surely have the capacity to insult and hurt, even if they are not libelous *per se*. On the other hand, like tainted ordinary preferences for goods and services, they can become relatively autonomous from their immoral origins and function for the discriminator no

⁶⁰ See *supra* text accompanying notes 40-42.

differently from aversions to spiders, snakes, and '40s big band music. Whether these preferences can be eradicated at little cost to the discriminator through cognitive psychotherapy is, I suppose, a matter of how deeply they are imbedded in the discriminator's basic personality and how autonomous they are from their progenitor judgments.

My conclusion is that some personal aversions—those that reflect nothing more than currently operative biases and/or inaccurate stereotyping, conscious or unconscious—are intrinsically immoral, whereas other personal aversions, if immoral at all, are only extrinsically immoral. Nevertheless, I must confess to being quite uncertain about even this rather hedged conclusion.

B. *Intrinsically Benign and Only Contingently Immoral Preferences*

1. Effects That Bear on the Morality of Intrinsically Benign Preferences

a. *Effects Caused by the Discrimination*

Discriminatory preferences that are not intrinsically immoral may nonetheless be extrinsically—contingently—immoral because of the effects on others of acting on those preferences. Although the most notable effects for assessing the extrinsic morality of discriminatory preferences were mentioned in Part II in my catalogue of such preferences, it is helpful to list them again here independently of the types of preferences that produce them. First, discriminating on the basis of some kinds of traits may often reinforce others' biases, immoral ideals, aversions, and inaccurate stereotypes, even if the discrimination itself is not of this nature. Use of race, for example, even where it is a rational proxy for an ordinary preference or a rational response to others' reactions, tends to feed residual racial biases, stereotypes, and aversions. These are serious negative effects that a discriminator should definitely consider, even though there is merit in the discriminator's likely claim that these effects are attributable to others' immoralities. It is, after all, morally reprehensible to leave an injured person in a high-crime area when he could be taken to safety with little hardship or risk, even though the danger he faces from abandonment is due to others' immoral tendencies.

A second set of negative social consequences consists of the psychic and motivational effects on members of the group defined by the disfavored trait. Members of the group whose sense of personal identity is very much bound up with such membership, by

observing so few other members in particular roles, may conclude erroneously that those roles are unrealistic aspirations and fail to develop talents that would be individually and socially rewarding. I say they may conclude this erroneously because most often, were they to develop the talents in question, the roles would be open to them. One exception to this is the use of a trait as a rational proxy, where what makes the proxy rational are the choices persons with the trait can be predicted to make in the present culture. Use of such proxies may tend to freeze the culture in its present state. For example, if because of accurate predictions of women's childbearing preferences, employers fail to promote them to certain jobs, women may tend to be reinforced in their childbearing preferences, thus freezing the cultural pattern, perhaps at great social cost in terms of lost productivity.⁶¹

One might argue that the discriminator should not be saddled with concern for social effects caused either by erroneous judgments about opportunities or by the actions of those countless others who contribute to making stereotypical judgments rational and thus to sustaining cultural patterns. I shall deal with the latter half of the claim when I deal with costs to the discriminator of forgoing use of certain intrinsically benign traits. As for requiring the discriminator to take account of others' misreadings of the situation he is helping to create, we frequently morally require people to take others as they are, with all of their tendencies to misjudge reality. For members of a group such as blacks, with their generally low socioeconomic status, poor education, and long history of oppression, it is quite natural to read the absence of blacks in a position as an indicator that efforts expended on attaining that position would be wasted, even if that is in fact untrue. Moreover, many of those to whom the "role model" argument applies are children, whose responsibility for misjudgments is diminished.

If we turn from motivational to psychic effects on members of dispreferred groups, we see effects that range from the resentment of "guilt by association" judgments represented by use of certain proxies to the loss of self-esteem suffered by those whose sense of well-being or most central sense of identity is inextricably bound to the fortunes of certain groups. An example of the former effect is that of a woman who may feel resentful that she is denied a position based on a quite rational prediction that she will choose eventually

⁶¹ See Kelman, *supra* note 21, at 1160-61, 1232-33.

to leave her current position to rear children because most women her age do so. An example of the latter effect is that of a black who may identify with the plight of blacks as a group, so that the low socioeconomic status of the group is felt by her as an injury she shares even though she is socioeconomically successful, or perhaps even as a crippling blow to her own self-esteem. These are the most common psychic harms produced by intrinsically benign preferences.

The case for requiring the discriminator to consider these psychic costs when deciding whether to discriminate on the basis of a particular trait is weaker than the case for requiring consideration of motivational effects. Although the resentment of "guilt by association" judgments is a natural resentment of the use of proxies that one "knows" do not predict one's own talents or behavior, proxy judgments are quite rational and inevitable. We cannot ask people to ignore at their peril statistical correlations between particular traits and unwanted behavior, even if the correlation is based on free choices.

The psychic harms from group identification are a product of both the discriminator's choice to use a trait to disfavor certain groups and the group member's choice to identify with the group defined by that trait and its fortunes. A manufacturer who wishes to employ only skilled widget makers may be quite indignant when told that she must consider that decision's costs to those who identify with the group defined as "those unskilled at widget making." Identification with such a group is not just strange and surprising. Rather, the identification looks like just another taste. It is also one that, if it wishes to be taken into account, should compete with others in the market. If people prefer that a disfavored group do better, it is they, rather than the discriminators in question, who should attempt to achieve this, perhaps by spending their own money on, for example, a voluntary private affirmative action plan.

Although the argument against considering psychic effects is plausible, a strong counterargument also exists. Although I am generally assuming a framework of rights, I have also pointed out in making this assumption that moral criticism of discrimination may lead us to revise our original estimate of where the boundaries established by these rights are located. If we assume in the case of psychic effects that the moral right lies with the discriminator, then of course it follows that those who identify with the dispreferred must compensate the discriminator if they wish to advance the dispreferred group. On the other hand, if these psychic costs cause

us to reconsider assigning the moral entitlement to the discriminator and to assign it instead to those with the psychic injuries, then the discriminator will be forced morally to internalize these costs of discriminating on the basis of the trait in question.

The appropriateness of assigning these psychic costs to the use of the trait in discriminatory decisions leads to a more general debate in normative theory. This debate turns on whether calculation of social welfare should include external preferences—meaning preferences about how other people should be treated or should fare—as opposed to personal preferences—preferences about one's own treatment and prospects.⁶² External preferences are preferences of particular people, and they are just as much preferences, and held at least as strongly, as purely personal preferences.⁶³ On the other hand, even if one rejects Ronald Dworkin's "double counting" criticism,⁶⁴ there are surely some bedeviling practical and perhaps theoretical difficulties in counting external preferences. If we count some people's preference that blacks as a group prosper in our assessment of the morality of using a trait that disproportionately adversely impacts on blacks as a group, do we also count other people's preferences that whites do better? Moreover, there will be a multitude of crosscutting group preferences and identifications. What helps blacks as a group might have a disproportionate adverse impact on other groups that are the subjects of psychic identifications and vicarious concerns, such as women as a group or Jews as a group. Should discriminators be morally saddled with the task of considering all these conflicting external preferences when deciding whether to act on an intrinsically benign preference?⁶⁵

I leave unresolved the question of how these psychic costs of group identifications bear on the morality of intrinsically benign discrimination. There remains one additional category of costs occasioned by such discrimination that deserves mention: the

⁶² See DWORKIN, *supra* note 8, at 232-38.

⁶³ See C. Edwin Baker, *Counting Preferences in Collective Choice Situations*, 25 UCLA L. REV. 381, 393 (1978); Gerald Dworkin, *Equal Respect and the Enforcement of Morality*, 7 SOC. PHIL. & POL'Y 180, 188-89 (1990); John H. Ely, *Professor Dworkin's External/Personal Preference Distinction*, 1983 DUKE L.J. 959, 961-64. For an excellent discussion of the role of preferences regarding others' preference satisfaction, welfare, status, etc., see Donald Hubin, *Non-tuism*, 21 CAN. J. PHIL. 441 (1991).

⁶⁴ See DWORKIN, *supra* note 8, at 234-38.

⁶⁵ Note the relation to the paradox of the Paretian liberal, which deals with problems associated with preferences regarding how others exercise their rights. See, e.g., Jonathan Riley, *Rights to Liberty in Purely Private Matters: Part I*, 5 ECON. & PHIL. 121, 133, 155-58 (1989).

category of civil strife and disharmony. For example, some members of some groups with which they identify feel such a stake in the fortunes of the group *qua* group that, to raise the group's status, they are prepared to engage in tactics that threaten social peace and harmony, which are themselves public goods. When private discrimination employs traits that thwart those groups' advancement *qua* groups, thus indirectly endangering social peace and harmony, it creates a negative externality by jeopardizing these public goods. Again, is this negative externality attributable to the discriminator for purposes of moral assessment of the discrimination, or is it attributable only to those group advocates who directly threaten social peace and harmony?

Finally, the social costs of particular kinds of discrimination are a function of how widespread those kinds of discrimination are. Harms associated with the absence of role models obviously occur only if the trait with which others identify is used quite frequently by private discriminators as the basis for exclusion of some from some benefit. One idiosyncratic use of a particular trait by a single discriminator is unlikely to affect the perception by members of the group defined by that trait of their general likelihood of obtaining positions and goods. For instance, if a particular employer wants his employees to have red hair, this is unlikely to affect brunettes' and blondes' perception of their life prospects and thus their motivation and development of talents. On the other hand, if many discriminators use the same trait to exclude, motivational and psychic effects are more likely to occur, especially if many people perceive their personal identity largely in terms of possession of that trait.⁶⁶ Although any individual discriminator's contribution to these costs is *de minimis*, each discriminator is contributing to the production of an external social cost in addition to the cost to the dispreferred person(s)—unlike the discriminator who idiosyncratically prefers redheads, a trait not specially preferred by many others.

⁶⁶ This is a favorite theme for those concerned about how groups are faring. See *supra* note 46 (citing articles focusing on the effect of civil rights initiatives on minorities, women, and the handicapped).

b. *Effects on the Discriminator of Forgoing the Discrimination*

In assessing the morality of discrimination that is not intrinsically immoral, the costs to the discriminator of forgoing the discrimination must also be taken into account in addition to the costs to the dispreferred, to those who identify with them, and to society in general. For some preferences, forbearance incurs high psychic costs. If an intrinsically benign preference is both central to the discriminator's self-definition and high in her hierarchy of values, then discrimination that is mandated by such a preference will be very costly for the discriminator to forgo. If the discrimination is based on a shallow preference—one forsworn with little psychic cost—or if the preference is merely a preference for the most efficient means of satisfying some further preference, then renouncing the preference will be less costly. If the trait defining the most efficient employees produces high social costs, whereas the trait that defines the slightly less efficient employees does not, choosing the latter trait over the former will be minimally costly to the discriminator but socially beneficial. Of course, losses of efficiency translate into monetary losses that ultimately translate into forgoing things that are intrinsically valuable to the discriminator. Thus, even if the discriminator is not deeply attached to the preference defining the most efficient means, she may be deeply attached to the preferences that the most efficient means enables her to realize. Nonetheless, some forms of discrimination, though rational, will be less costly to forgo than others.

Another factor to consider in assessing the costs to discriminators of forgoing certain forms of discrimination is whether the costs can be eliminated or reduced once collective action problems are eliminated. For instance, many employers in choosing employees would be quite willing to overlook immoral consumer reactions if assured that their competitors would also do so. The consumers, if faced with the choice between forgoing the service and dealing with employees against whom they were biased, might all choose the latter course of action. In such a case, no employer would lose any profits if all employers joined in or were forced into refraining from catering to consumer biases.

In the cases where the discriminator is considering forgoing use of traits defining efficient employees or forgoing use of efficient proxy traits, solving collective action problems might eliminate *some* of the costs to employers of forgoing the use of such traits, such as losses to competitors who continue to discriminate on the bases of

those traits. It is not likely to eliminate all of the costs, however, since forgoing the use of the efficient traits will raise the price of the good or service and probably reduce the demand. Realistically, asking discriminators to forgo efficient traits, proxy or non-proxy, will impose costs on them that can only be partially passed on to others even if collective action problems are eliminated.⁶⁷

C. *Immutability of Traits*

How should a trait's immutability or relative costliness to change bear on the morality of discrimination on the basis of such a trait? If discrimination is based on judgments of relative moral worth, then, of course, we good Kantians are likely to reject such judgments if they are based on immutable characteristics. Moral worth must be based on moral choices, not on physical characteristics or even character traits to the extent that such traits are not just proxies for the prior moral choices that formed them.

If we shift the focus from biases to other forms of discrimination, however, the relative immutability of a trait seems to drop out as an independent factor. People will rarely deeply identify with traits that are easily changed, especially if we factor psychic costs into the determination of how changeable a trait is. One's religious identification is a trait that is easily changeable in one sense, but not in the sense that I think is relevant. Gender is now mutable in the sense that sex-change operations are available. But the monetary expense of a sex-change operation—which can be translated into the psychic costs of forgoing other goods that the money could buy—as well as the psychic costs that most people would suffer from undergoing such a procedure, make gender immutable in any relevant sense. Indeed, the psychic costs and basic opportunity costs associated with changing a trait with which one identifies define the relevant immutability of a trait. Those psychic costs, however, have already been covered in the previous discussion.

⁶⁷ See Michael H. Gottesman, *Twelve Topics to Consider Before Opting for Racial Quotas*, 79 GEO. L.J. 1737, 1746-48 (1991).

D. Immoral Discrimination and Moral Rights

1. The Moral Rights of the Discriminator

The previous discussion has focused on the factors relevant to determining whether the discriminator is acting immorally in preferring and disfavoring certain traits. Determining that the discrimination is immoral, however, may not answer the question whether the discriminator is acting within his or her moral rights. One may have the moral right to act immorally in some circumstances. The value of autonomy warrants ascribing to agents a moral space—a space protected by moral rights—in which the agent may exercise that autonomy in a variety of ways, including ways that are intrinsically or extrinsically immoral. For example, a person who in choosing a spouse or an intimate companion excludes members of a particular race solely because of a bias may be acting within her moral rights even if she is acting immorally. Likewise, a person who, based on religious convictions that women should stay at home or that blacks are inferior, refuses to let women preach in or to let blacks join his church, may be acting within his rights even if immorally.

In saying that one may have a moral right to engage in intrinsically immoral discrimination, I am adopting an orthodox though not uncontroverted view of the moral universe.⁶⁸ Some thoroughgoing act utilitarians might, of course, deny the gap between moral rights and morally right actions. Most act utilitarians, however, would probably, for strategic reasons, grant moral rights to act immorally if we construe moral rights as morally warranted claims against coercion.⁶⁹ So, too, would even those nonutilitarians such as Joseph Raz, who deny any moral value to the exercise of autonomy when it is expressed in immoral choices.⁷⁰ Raz argues that although the value of autonomy requires for its realization a range of alternative projects and lifestyles from which to choose, it does not require that any of those alternatives be immoral ones.⁷¹ In fact, Raz argues, autonomy's moral value is fully cashed out in the moral value of the objects of choice, so that the autonomy to opt for what

⁶⁸ See Jeremy Waldron, *A Right to Do Wrong*, 92 ETHICS 21, 34-37 (1981).

⁶⁹ See, e.g., JOHN GRAY, *MILL ON LIBERTY: A DEFENCE* 48-57 (1983).

⁷⁰ See RAZ, *supra* note 4 at 381; see also Donald H. Regan, *Authority and Value: Reflections on Raz's MORALITY OF FREEDOM*, 62 S. CAL. L. REV. 995, 1082-85 (1989).

⁷¹ See RAZ, *supra* note 4, at 410-12.

is immoral can itself have no positive moral value.⁷² But despite his position that there is no moral value in autonomy exercised immorally, Raz rejects the use of coercion—legal sanctions—to deter immoral choices simply because they are immoral.⁷³ Even Raz would probably grant unqualified legal rights to bigots to choose their spouses or co-religionists or to dispose of some forms of personal property in bigoted ways.

2. The Moral Rights of Victims of Immoral Discrimination

The costs to others imposed by immoral discrimination are not only part of the calculus by which its (extrinsic) immorality is determined; they may also bear on the question of the victims' moral rights to be free of the discrimination. I assumed as a hypothetical background for this inquiry that victims' just claims were fully honored, and I stated that in general I would be attempting to ascertain the morality of discrimination on the basis of that assumption; however, I also stated that the boundary line separating claims of denial of one's just share from other moral objections to discrimination might be fixed, at least in part, by a determination that the discrimination is immoral. For example, if we conclude that an employer's immoral discrimination in selecting employees does not fall within the realm of her moral rights, then we might conclude that the rights of prospective employees include freedom from at least some of this immoral discrimination. Put differently, justice may include a right to be free of seriously immoral forms of discrimination in employment. On the other hand, there will be some, perhaps many, types of discrimination that, though immoral vis-à-vis the victims, and though not within the realm of the discriminator's moral rights, are not violative of the victims' moral rights.

3. Immoral Discrimination and Moral Rights: A Summary

The way in which the morality of various types of discrimination affects the boundaries of the discriminators' and the victims' moral rights ultimately depends upon the background moral theory from which those rights are derived. For example, a derivation from a thoroughly consequentialist theory might be more sensitive to the

⁷² See *id.* at 373-81.

⁷³ See *id.* at 403-04, 418-19.

effects of various types of discrimination than a derivation from a non-consequentialist moral theory.

The map of the moral universe that emerges from the previous discussion looks like this. Some immoral discrimination (and lots of morally permissible discrimination) falls within the discriminator's moral rights. Some immoral discrimination violates its victims' moral rights. (No morally permissible discrimination could violate others' moral rights since it would for that reason not be morally permissible.) If additionally one believes that moral rights do not cover the realm of all possible actions, then there is a middle realm, of uncertain size, between the moral rights of the discriminator and those of the victims. In that middle range are located some, and perhaps many, immoral forms of discrimination, along with perhaps some morally permissible forms.

E. *The Law and Immoral Discrimination*

The law, properly employed, will attempt first to secure moral rights, both those of the discriminators and those of the victims of discrimination. Because the law is a blunt instrument, however, and its necessary formality and enforcement imperfections will cause it to deviate from precisely tracking the subtle boundaries of moral rights, it will not and should not be designed to produce total enforcement of moral rights. In some cases, because the space between discriminators' and victims' moral rights is so small or is nonexistent, total enforcement of victims' moral rights will jeopardize discriminators' moral rights, and vice versa. In other cases, the costs of total legal enforcement—monetary costs and/or costs of legal intrusiveness—will jeopardize other important values, including moral rights. The law then should not concern itself with all possible violations of moral rights, but should instead select for prohibition those broad types of discrimination that are most likely to be immoral, intrinsically or extrinsically, that are either likely to violate victims' rights or to cause a great amount of social harm, that are least likely to be matters of discriminators' moral rights, and that are least costly to detect and establish in court.

How well does the American law of private discrimination, at least on the federal level,⁷⁴ satisfy these criteria for legal prohibition?

⁷⁴ My focus here is on the Civil Rights Act of 1964, 42 U.S.C. §§ 2000a-2000e (1988); the Fair Housing Act of 1968, *id.* §§ 3601-31; the Equal Pay Act of 1963, 29 U.S.C. § 206(d) (1988); the Age Discrimination in Employment Act of 1967, *id.* § 623;

Some of the prohibited traits on which to base discrimination—race, sex, religion, and national origin—are traits a high percentage of use of which will reflect biases, unjustified ideologies, inaccurate stereotypes (usually bias-driven), shallow personal aversions (usually bias-driven), and reaction qualifications where the reactions are immoral. Some of the traits, specifically race and sex, also identify types of discrimination that have been widespread and socially costly. On the other hand, this seems to be less true—though not completely untrue—of the traits of old age and handicap.

An additional point in favor of the federal anti-discrimination laws—leaving aside the 1866 civil rights laws embodied in §§ 1981 and 1982 of the United States Code⁷⁵—is that they limit their prohibitions to contexts where the moral rights of discriminators are least likely to be implicated. Discrimination is prohibited by employers or landlords in businesses or apartments that are large enough not to involve personal relationships or personal aversions as major concerns in hiring or renting.⁷⁶ Discrimination by small employers and landlords—and by employees and consumers in all contexts⁷⁷—is not prohibited. There are also exemptions for religious organizations.⁷⁸ The scope and limitation of the laws on discrimination in public accommodation or in the use of federal funds show the same solicitude for the discriminators' moral rights.⁷⁹ There is, of course, no ban on discrimination in the context of personal relations.

The bona fide occupational qualification ("BFOQ") component of federal employment discrimination law also fits the prescription for legal intervention rather well.⁸⁰ Sex or religion will only rarely

the Rehabilitation Act of 1973, *id.* §§ 793-94; and the Civil Rights Act of 1866, 42 U.S.C. §§ 1981-82 (1988).

⁷⁵ See 42 U.S.C. §§ 1981-82 (1988).

⁷⁶ See *id.* §§ 2000e(b), § 3603(b)(2).

⁷⁷ Interestingly, not even the most sweeping civil rights laws, such as §§ 1981 and 1982, which have been construed (erroneously, in my opinion) by the Supreme Court to reach private discrimination, see *City of Memphis v. Greene*, 451 U.S. 100, 120-24 (1981), have ever been applied to discrimination *against* employers by prospective employees, or to discrimination *against* sellers by prospective buyers. See 42 U.S.C. §§ 1981-82 (1988); Lawrence A. Alexander & Maimon Schwarzschild, *Consumer Boycotts and Freedom of Association: A Comment on a Recently Proposed Theory*, 22 SAN DIEGO L. REV. 555, 560-62 (1985).

⁷⁸ See 42 U.S.C. §§ 2000e-2(e), 3607 (1988).

⁷⁹ See, e.g., *id.* § 2000a(e) (stating that private clubs are not subject to restrictions on discrimination in public accommodations).

⁸⁰ Employers may use religion, sex, or national origin as a factor in their decisions "to hire and employ . . . in those certain instances where religion, sex, or national

qualify as BFOQs; therefore, allowing discrimination based thereon when they are BFOQs will not produce major external social costs. Moreover, BFOQ law becomes confusing exactly where one would expect—where what purports to make sex and religion BFOQs are consumer reactions whose moral status is ambiguous. For example, the courts have treated airlines' preferences for female cabin attendants as violative of the law without regard to the flying public's clear preference for females in that role.⁸¹ On the other hand, hiring females as strippers or as models for dresses—where being female is a BFOQ only given a public preference for sex-differentiated striptease shows and sex-differentiated clothing—has been deemed permissible.⁸²

Use of the forbidden grounds as rational proxies, in contrast to their use as BFOQs, is forbidden in the employment context. Although sometimes race and sex will in fact be rational proxies for an employer to use, there are reasons for making their use per se illegal. First, even their use as rational proxies may be extrinsically immoral because of the social effects of such use. Second, many claims that the use of such grounds as proxies is rational will be covers for the expression of immoral biases or aversions. Third, even if their use as proxies is not immoral, forbidding their use in the contexts in which the employment discrimination law applies ("largish" firms) is quite unlikely to infringe employers' moral rights. Finally, a per se ban on their use as proxies may be quite preferable to an attempt to adjudicate the legitimacy of each use on a case by case basis.

On the other hand, legal suppression of quite cost-effective proxies will produce a strong incentive to find close substitutes for the forbidden proxies.⁸³ If such substitutes can be found, the costs of forbidding the original proxies will be largely averted along with the social costs of allowing their use. If there are no close substitutes for the forbidden proxies, however, the prohibition will result in imposition of substantial costs on the business in question, their

origin is a bona fide occupational qualification reasonably necessary to the normal operation . . . [of the employer's] business or enterprise." *Id.* § 2000e-2(c)(1). This exception is not permitted when the discrimination is based on race or color.

⁸¹ See *Diaz v. Pan Am. World Airways*, 442 F.2d 385, 387 (5th Cir.), *cert. denied*, 404 U.S. 950 (1971).

⁸² See ARTHUR LARSON & LEX K. LARSON, *EMPLOYMENT DISCRIMINATION* § 15.10 (1992).

⁸³ See *Williams*, *supra* note 46, at 1780-82.

employees, and their customers, and may even result in harm to the prohibition's intended beneficiaries.⁸⁴

In the realm of private insurance, use of race, sex, etc., as proxies is not currently forbidden, at least at the federal level, when the insurance is offered by an insurance company and not by the employer.⁸⁵ Insurance, of course, is by its very nature built upon proxies. But we could require that risks not be differentiated along certain lines even if such differentiation is justified by resource allocation considerations.

With an important qualification, ordinary and tainted preferences that result in a disparate impact along racial, sexual, and other lines—and no preference that I can imagine, except a preference for selection by random procedures, will fail to produce a disparate impact on either blacks or whites, or on either females or males, and so forth—are not prohibited, even if the disparate impact has, or in any event is alleged to have, major social consequences. Employment discrimination law took a substantial step toward outlawing all preferences for employees that produce disparate impact along the named axes in the 1971 case *Griggs v. Duke Power Co.*,⁸⁶ in which the Supreme Court held that an employment practice disproportionately and adversely affecting blacks' employment opportunities must

⁸⁴ See RICHARD A. EPSTEIN, *FORBIDDEN GROUNDS: THE CASE AGAINST EMPLOYMENT DISCRIMINATION LAWS* 242-66, 258 (1992) (arguing that "the application of [Title VII] to private discrimination in competitive markets generates not only administrative costs but allocative losses as well. . . . The real private gains, shared by blacks and whites alike . . . are diminished when private choices are subjected to systematic state scrutiny.").

To say there are costs associated with prohibiting the use of certain cost-effective proxies is not to say of course that such a prohibition is unjustifiable. The associated social costs of using the proxies may be greater than the costs of their prohibition or may be less fairly distributed. Moreover, the costs may represent a tax to effect a crude form of reparations to those groups benefitted by the prohibition, reparations that may be justified in our society but not in the more ideal society that I am assuming in this article. I owe this last point to Ken Simons. See *supra* note 38. I am skeptical of the tenability of such a reparations justification. See *supra* note 51. I am equally skeptical that the costs of prohibiting cost-effective proxies can be justified as a tax in furtherance of a more just distribution of wealth; there are undoubtedly a multitude of less crude, less costly—not just in general, but also to many of the poor—and less racially-skewed alternatives.

⁸⁵ For an excellent discussion of employer-provided insurance and the federal antidiscrimination laws, see EPSTEIN, *supra* note 84, at 340-49. See also Martin J. Katz, *Insurance and the Limits of Rational Discrimination*, 8 YALE L. & POL'Y REV. 436, 437-38 (1990) (discussing state regulation of insurance companies' choices of proxies and noting that "no state has instituted a comprehensive prohibition of 'impact' rational discrimination").

⁸⁶ 401 U.S. 424, 431 (1971).

be justified by business necessity. The case could be viewed as establishing an evidentiary and procedural framework for smoking out covert uses of forbidden criteria⁸⁷—the absence of business necessity resulting in a conclusive presumption of such uses—and it was applied only when the adverse impact fell on blacks rather than whites or on females rather than males.⁸⁸ (Otherwise, every employment practice would be vulnerable to the requirement that it be justified by business necessity.) The Supreme Court retreated from *Griggs* to some extent in *Wards Cove Packing Co. v. Atonio*,⁸⁹ though Congress has attempted to reinstate *Griggs* in a new civil rights bill signed by President Bush.⁹⁰ The underlying rationale for forbidding preferences for employees that produce adverse disparate impact on certain groups remains unclear when the preferences do not reflect biases and are otherwise cost-justified. Perhaps the rationale is the same as that behind forbidding the use of certain traits as proxies, even if such use is currently cost-justified.⁹¹ Interestingly, the law distinguishes between search costs—which the *Griggs* approach requires the employer to incur by precluding cost-justified employment criteria—and training costs, which are not required by the *Griggs* approach.⁹² In other words,

⁸⁷ See CHARLES FRIED, ORDER AND LAW 93-95 (1991).

⁸⁸ Interestingly, the Court refused to extend *Griggs* to practices that adversely affected groups of foreign national origin. See *Espinoza v. Farah Mfg. Co.*, 414 U.S. 86, 94-95 (1973).

There are two aspects of the *Griggs* rule that suggest that its rationale is not one of smoking out covert uses of forbidden criteria deemed unfair to excluded individuals, but is instead one of legally establishing a presumption of proportional representation of certain groups *qua* groups in the workplace. First, there is no cause of action available to, for example, an individual black excluded by criteria that are arbitrary relative to business goals but that do not disproportionately disadvantage blacks as a group, even if criteria that would pass the business necessity test would have produced a workforce in which blacks were "overrepresented." Second, there is a cause of action available to any black excluded by criteria that fail to satisfy the *Griggs* rule even if it is likely that criteria that would have passed the business necessity test would also have excluded more blacks. On the other hand, this proportional representation interpretation of the *Griggs* rule is to some extent countered by the fact that the courts have not applied *Griggs* where whites have been the disproportionate impact victims of employment criteria.

⁸⁹ 490 U.S. 642, 650-52 (1989).

⁹⁰ See Civil Rights Act of 1991, Pub. L. No. 102-166, § 107(a), 105 Stat. 1071, 1075 (1991). For a discussion of the "critical transition" from disparate treatment to disparate impact via *Griggs*, *Wards Cove*, and the Civil Rights Acts, see EPSTEIN, *supra* note 84, at 182-204.

⁹¹ See *supra* text accompanying notes 82-84.

⁹² See Gottesman, *supra* note 67, at 1750-52; Kelman, *supra* note 21, at 1162, 1202-03.

if a protected group is as talented at a job as other groups, but identifying its talented members is more costly than identifying the talented members of other groups, the employer must bear the costs. If a protected group is less talented at the job, however, the employer need not train its members, even if the costs of training are less than the identification costs.

Finally, it is permissible for an employer to discriminate on the basis of the otherwise forbidden criteria if the employer is implementing an ideal of an integrated workplace. In short, racial and other private affirmative action plans are permitted.⁹³

F. Three Current Issues: Comparable Worth, Disparate Impact, and Affirmative Action

It will be instructive to conclude my investigation of the morality of discrimination by focusing on three quite controversial issues regarding the legal response to discrimination. The first issue is whether it should be considered illegal sex discrimination for an employer to have different wage scales for different job categories where the jobs are of equal "comparable worth," and where the lower paid job is one predominantly held by women and the higher paid job one predominantly held by men (the "comparable worth" issue). The second issue is whether discrimination on the basis of those traits that lead to a disproportionate exclusion of minorities and women should be treated either as illegal per se or illegal in the absence of a compelling justification (the "disparate impact" issue). The third issue is whether private affirmative action employment plans that favor women and minorities are legal and/or morally warranted (the "affirmative action" issue). I shall take these issues up in turn.

Before doing so, however, I should comment on the limitations of my prior analysis of discrimination in assessing the real world problems of comparable worth, disparate impact, and affirmative action. My prior analysis assumes a just society. The issues that I deal with here are issues that confront our society as it is, which is surely not perfectly just by anyone's standards. Moreover, to the extent that these issues arise in debates over legal policy, straightforward application of my analysis is hindered by the fact that discrimi-

⁹³ See *United Steelworkers v. Weber*, 443 U.S. 193, 208 (1979). This is consistent with forbidding certain cost-justified proxies, as both represent employers' bearing bottom-line losses to advance certain groups' prospects. See Sunstein, *supra* note 21, at 28-29.

nation need not be immoral, intrinsically or extrinsically, to be justifiably and legally prohibited. So long as the prohibition does not infringe discriminators' moral rights and is otherwise good social policy, the discrimination it bars may be morally permissible.

Although these limitations, coupled with the relevance of factual determinations that are outside the domain of this article, preclude my making strong claims about comparable worth, disparate impact, and affirmative action, they do not render my analysis completely beside the point. If in our actual social situation a sound policy can be predicted to be very costly along some dimensions, then its benefits in terms of furthering social justice must outweigh those costs—costs that can also be converted into social justice currency—and it must be the least costly means available for achieving those benefits. I am skeptical that proposals for achieving comparable worth and proposals for eliminating disparate impact (including affirmative action plans) meet these conditions. Even if I am wrong in my assessment of these issues, however, the validity of my general framework for the moral analysis of private discrimination is unaffected.

1. The Comparable Worth Issue

Most job categories are not fully integrated in terms of gender, so that there are many jobs predominantly held by males and many jobs predominantly held by females. This situation can be explained by a number of factors, including: 1) some present-day discrimination against women, usually illegal and covert; 2) some inborn differences between men and women in terms of certain skills, strengths, and aptitudes; 3) some present-day disabilities caused by past discrimination in education and training opportunities; 4) some real and untainted differences in preferences; and 5) some differences in preferences traceable to ideologies of gender role ("women's work" and "men's work") and to past discrimination ("don't aspire to jobs you'll be denied"). When two jobs require "comparable" skill levels and training, and one is predominantly held by males and the other predominantly by females, the predominantly male job will sometimes receive dramatically higher remuneration. This situation is what concerns the proponents of comparable worth theories of illegal discrimination. They argue that unequal pay for jobs of "comparable worth" should be deemed illegal gender discrimination, at least where women hold the lower paid job.

In any comparable worth situation it is always possible that there is an intrinsically immoral gender preference being covertly expressed by an employer. Or the employer may be surreptitiously using gender as a proxy or as a reaction qualification under conditions that make such use extrinsically immoral. In either case, the covert use of gender will likely be illegal under current antidiscrimination laws without the necessity for treating the practice of unequal pay for jobs of comparable worth as itself illegal gender discrimination. Thus, unequal pay for jobs of comparable worth may be evidence of ordinary forms of illegal gender discrimination even if it is not illegal *per se*.

In cases where the employer is merely responding to the supply of and demand for workers, or to customer demand for the goods and services she produces, is there any reason to deem the employer's actions immoral and subject to legal prohibition?

The basic cause for unequal pay for jobs of "comparable worth" is that women and men are particularly attracted to certain jobs and shun other, higher-paying jobs for which they are otherwise qualified. They may prefer to be nurses rather than truck drivers, school teachers rather than plumbers, and so on. With many women preferring the same jobs, the wage for that job will be driven down by market forces. If the employer is compelled by moral and/or legal condemnation to pay more than the market-clearing wage, distortions will reverberate throughout the economy. Moreover, the employer is limited in what she can pay employees by the consumer demand for her products. I might choose to undergo years of training in order to be able to make sixteenth century lutes. By "comparable worth" reasoning, I should be paid the salary of a neurosurgeon. But what if there is no demand for my lutes comparable to the demand for neurosurgery?

Assuming that it is not unjust to have a market economy, which includes perforce a market in labor, and assuming that differentials in pay that reflect the supply of and demand for workers are not generally immoral—even when those differentials exist for jobs of "comparable worth"—it is difficult to see the case for moral and legal condemnation of such differentials in the context of gender.⁹⁴ Is the case for comparable worth bolstered by the fact that the differences in job preferences between males and females are the

⁹⁴ Unlike the disparate impact concern for group equality that I discuss in the next section, the comparable worth claim is limited to equality of only those males and females with comparable skills and training levels.

product of tainted preferences, preferences traceable to ideologies of gender roles or to past discrimination against women?⁹⁵ I have already said that discriminating on the basis of tainted preferences is not intrinsically immoral.⁹⁶ Are the negative social effects of unequal pay for jobs of comparable worth such that we should conclude that such discrimination is extrinsically immoral when traceable to tainted preferences? Although there are social harms that occur when, for whatever reasons, individuals do not seek the most socially valued (in the sense of highest paid) jobs for which they are qualified, those harms are best redressed through actions that do not radically distort the workings of the market. (Indeed, those harms are usually offset by the utility gain to those individuals that induced them to take the less socially valued job, assuming they were not misinformed or the victims of false consciousness.) The harms that would be suffered by the employer were she compelled to pay above-market wages—harms that reflect the general social harms of market distortion—surely seem greater than any social harm caused by unequal pay for jobs of “comparable worth.” In any event, requiring the employer to pay more than the market rate, rather than redressing the social harm of individual underachievement, actually exacerbates that social harm by providing incentives to go into jobs for which there is already an oversupply of qualified labor.⁹⁷ Ultimately, comparable worth’s logical conclusion is the elimination of the free market and its replacement by a command economy.

⁹⁵ See Littleton, *supra* note 33, at 1296-1301. Littleton argues that her favored conception of equality—“equality as acceptance”—requires society to treat as equally valued men’s and women’s “genderized” (read: tainted) values and inclinations and to devote equal resources to male and female occupational choices. Littleton omits discussion of the mechanics and market effects of her proposal and leaves even those politically close to her unpersuaded. See, e.g., Kelman, *supra* note 21, at 1191-92 & n.90, 1193 n.93 (disagreeing with Littleton’s affirmative distributive principle that distinct social subgroups have claims to equal per capita income); Mark Kelman, *Emerging Centrist Liberalism*, 43 FLA. L. REV. 417, 440-42 (1991) (disagreeing with Littleton’s case for comparable worth).

⁹⁶ See *supra* text accompanying notes 59-60.

⁹⁷ See Richard A. Posner, *An Economic Analysis of Sex Discrimination Laws*, 56 U. CHI. L. REV. 1311, 1329-31 (1989).

2. The Disparate Impact Issue

Currently, it is not illegal discrimination to prefer traits that are disproportionately distributed among races and ethnic groups or between the genders. The disparate impact caused by such preferences may be evidence that the preferences are really covers for racial, ethnic, or gender preferences that are themselves illegal biases, proxies, aversions, or reaction qualifications. Much of the controversy over when the burden of proof in employment discrimination cases shifts from the employee to the employer, and what the employer must prove when the burden is shifted, is about the role of disparate impact as evidence of illegal discrimination.⁹⁸

There are, however, those who want disparate impact on minority races and ethnic groups as well as on women to be treated as illegal discrimination *per se*, at least in the absence of an employer justification that is compelling, and not just sufficient to negate the inference of covert employer use of race, ethnicity, or gender. Accordingly, even if the employer really does prefer the trait in question for otherwise legitimate business reasons, use of the trait should be prohibited, or so it is argued, if such use has a disparate negative impact on certain groups, unless the employer can point to disastrous consequences that use of a more evenly distributed trait would produce. Furthermore, because there are a multitude of traits employers may prefer that will be more commonly found in nonminority males than in others, recognition of the disparate impact cause of action would have quite far-reaching consequences.

Is it immoral for an employer to prefer a trait that is disproportionately uncommon in minority groups and women? It is surely not obvious why this should be so, and proponents of disparate impact illegality are usually quite opaque about their underlying rationale. But there are two grounds for this position that seem most plausible. First, there are negative social effects associated with some groups' low socioeconomic status, and this status in turn is the product of the preferences for traits that have a disparate impact. Second, even if these negative effects are insufficient in themselves to deem the preferences (extrinsically) immoral, when these effects are coupled with the historical explanation for why the traits are unevenly

⁹⁸ At least this was so after *Wards Cove Packing Co. v. Atonio*, 490 U.S. 642, 650-52 (1989), which eliminated the requirement from *Griggs v. Duke Power Co.*, 401 U.S. 424, 431 (1971), that employers justify disparate impact by proving business necessity. See *supra* text accompanying notes 86-92.

distributed—an explanation in which immoral discrimination is central—the immorality of disparate-impact-producing preferences emerges.⁹⁹

For reasons I have already given in my discussion of the morality of discriminating on the basis of ordinary preferences,¹⁰⁰ I find the case for disparate impact illegality quite unpersuasive. First, if the disparate impact is traceable to the effects of past immoral and/or illegal discrimination on the distribution of preferred traits, that calls for reparations perhaps as a matter of justice, but not for ignoring the traits people currently possess. Up until this section, I have been assuming a framework of justice, one where all valid claims for reparations have been honored. Suppose we now drop this assumption and ask the following question: If one lacks a preferred trait that one would have possessed were one not the victim of wrongful discrimination, may one demand in addition to compensation from one's victimizers that others overlook the absence of the trait? Surely not. If I was a promising neurosurgeon before my hands were mangled by a drunk driver, I cannot ask others to ignore my lack of dexterity in their choice of neurosurgeons. If I was a beauty before an assault left my face scarred, I cannot ask others to pick me to model. Preferences for scarce traits do not become immoral merely because the scarcity is in part due to the immoralities of others, even when just reparations have not been fully paid.

That leaves us to consider the disparate impact itself without regard to its causes. That is, is the fact that a preference will produce disparate group impact along certain axes a cause for moral concern in the unjust real world?

It is undeniably the case, as I earlier pointed out, that many people, both within and without particular groups, care greatly about how those groups *qua* groups are faring. Given alternative worlds in which there are 100,000 doctors and 100,000 manual laborers, they prefer the world in which 50,000 of each position are black and 50,000 of each position are white to the world in which 80,000 doctors are white and 80,000 laborers are black. This preference holds, moreover, even if we eliminate all inferences from disparate group impact to current racial discrimination.

⁹⁹ See Strauss, *supra* note 21, at 1654 n.58.

¹⁰⁰ See *supra* text accompanying notes 46-54.

Should a discriminator be morally and legally compelled to consider the effect of his preferences on others' group preferences? As I pointed out in discussing ordinary preferences, in a just society, the affirmative case is quite weak.¹⁰¹ It is perhaps strengthened somewhat by pointing to a political divisiveness effect beyond mere disappointment of others' group preferences. And it might be strengthened even further by plausible claims that disparate impact reduces aspirations caused by absence of role models within one's group, or by plausible claims that disparate impact reinforces biases and inaccurate negative stereotypes. Moreover, if we drop the assumption of a just society, the case is further strengthened if disparate impact disadvantages groups that correlate positively with the group of unjustly disadvantaged.

There may, therefore, be some ordinary preferences that are extrinsically immoral because they are shallow and easily forgone, and because, as a result of disparate impact, they frustrate preferences for group advancement, contribute to political strife, reinforce biases and stereotypes, sap motivation, and impact disproportionately individuals who are unjustly disadvantaged. Nevertheless, the case seems quite shaky for blanket legal condemnation of all the ordinary preferences that contribute to inequalities of socioeconomic status of groups defined along certain axes, and it is untenable if there are plausible and less costly alternatives to forbidding discrimination that produces disparate impact.¹⁰²

¹⁰¹ See *supra* text accompanying notes 62-65.

¹⁰² Nonetheless, the number of commentators calling for such legal condemnation is disconcertingly large. For representative examples, see Eichner, *supra* note 46, at 1416-17 (arguing that job requirements disproportionately excluding women cannot be justified by the costs of alternatives that would exclude women less, and that such costs must be weighed against the value of equality—an idea I find difficult to grasp except as part of a more general position in favor of equalizing wealth across the board); Littleton, *supra* note 33, at 1284-85 (advocating equal treatment of all "genderized" occupational choices in terms of resources, thereby virtually calling for gender group equality of income); McCluskey, *supra* note 46, at 878-80 (treating inequality among certain groups *qua* groups as wrong per se); see also Becker, *supra* note 46, at 207-08.

Of course, if one drops the assumption of a basically just framework, worrying about the equality effects of various preferences is more understandable, though not the equality effects among groups *qua* groups as opposed to the equality effects among individuals *qua* individuals. See, e.g., BRUCE A. ACKERMAN, *SOCIAL JUSTICE IN THE LIBERAL STATE* 18 (1980); RAWLS, *supra* note 8, at 504-12; Ronald Dworkin, *What Is Equality? Part 2: Equality of Resources*, 10 PHIL. & PUB. AFF. 283, 283-90 (1981); see also Daniel Wikler, *Paternalism and the Mildly Retarded*, 8 PHIL. & PUB. AFF. 377, 392 (1979).

This discussion of disparate impact may be the appropriate place for bringing together various ideas that have been mentioned at different points in the body of this article. In my background assumptions I stated a general presumption in favor of macro policies and against regulation of private preferences for effecting social justice.¹⁰³ At another point in this article I stated a presumption against moral positions that exact great individual sacrifices for the general good.¹⁰⁴ At still another point I raised the possibility that ordinary preferences that produce negative social consequences might be immoral only if those with the preferences were cheaper cost avoiders than those whose responding preferences and reactions produced the negative social consequences.¹⁰⁵

These points are all related, and they all bear on why ordinary preferences that produce disparate adverse impact on particular groups will infrequently be immoral. If a particular ordinary preference produces disparate impact, and there are plenty of close substitutes available that would not, we might believe that those with the preference should give it up and switch to the substitutes. Such situations, however, are rarely ones where the disparate impact produces major social harm. Disparate impact produces major social harm where countless ordinary preferences and their substitutes will produce the disparate impact. For example, where the production of Coca-Cola requires skills not frequently found among the members of a particular social group that is concerned with its status as a group, but the production of Pepsi and RC Cola does not require such unevenly distributed skills, little would be lost and perhaps much would be gained by morally forbidding Coca-Cola production. Disparate impact, however, gives rise to social costs only when it is the product of a large number of different preferences, not just preferences for Coca Cola, but preferences for Pepsi, RC Cola, and a multitude of other goods and services as well. It is when members of the same self-conscious group disproportionately lack the skills necessary to satisfy all sorts of preferences, not just very specific ones, that the negative social effects of disparate impact are produced.

If this is the case, then those with ordinary preferences are unlikely to be the cheapest cost avoiders, because, to avoid the costs, they would have to give up, not just particular preferences for which

¹⁰³ See *supra* text accompanying note 3.

¹⁰⁴ See *supra* text accompanying note 23.

¹⁰⁵ See *supra* text accompanying notes 49, 62-67.

there were available close substitutes, but a whole way of life. It is no doubt true that our general preference for a society with a highly advanced technology has been hard on the self-esteem of those with skills most suitable for a primitive agrarian society.¹⁰⁶ Additionally, it is quite possible that a general preference for a primitive agrarian society would result in a more egalitarian distribution of esteemed traits. Nonetheless, few of us would condemn the preferences for technology and the benefits it brings merely because we are endowed unequally with the traits necessary to produce those benefits.

When disparate impact and its negative effects are traceable to countless ordinary preferences, we are unlikely to deem those with the preferences to be the cheapest cost avoiders. To cure the negative social effects we are more likely to turn our attention to macro policies—education, job training, redistribution of wealth—the burdens of which fall much more uniformly across society, unlike condemnations of particular ordinary preferences.¹⁰⁷ And those whose group identifications and sympathies generate the social costs of ordinary preferences will begin to look much more like the cheapest cost avoiders than will those with the ordinary preferences. Moreover, if we focus on helping the group defined as those individuals who lack just shares of social benefits, and not on racial, ethnic, gender, and similar social groups, macro policies look obviously superior to attempting to change preferences and thereby affect the share of resources various talents can command.¹⁰⁸

¹⁰⁶ Cf. Wikler, *supra* note 102, at 386-90 (discussing the social preference for a technologically advanced way of life as it affects the ability of the mildly retarded to live without being subjected to legal guardianship).

¹⁰⁷ Mark Kelman notes the similarity between the costs of identifying those with qualifications for a particular job and the costs of training those who would be qualified if trained. From the employer's standpoint, costs are costs. Kelman concludes that although both costs are utilitarian concerns, they are both offset by a deontological claim that those with equal pre-training talent should have equal chances in employment. See Kelman, *supra* note 21, at 1231-33, 1245-47. I cannot see the deontological force of a claim based on raw talent, though others besides Kelman do. Compare RAWLS, *supra* note 8, at 84, 87, 302-03 (endorsing a fair equality of opportunity principle for those of equal talent) with Larry A. Alexander, *Fair Equality of Opportunity: John Rawls' (Best) Forgotten Principle*, 11 PHIL. RES. ARCHIVES 197, 206 (1986) (arguing that Rawls fails to provide a good explanation of why his fair equality of opportunity principle would be favored independent of and prior to his general difference principle).

¹⁰⁸ I suspect that the notion of cheapest cost avoider is central to the assessment of discriminatory preferences in ways beyond those just mentioned. For example, concerns with "immutable" traits can, I think, be translated into cheapest cost avoider concerns. So, too, perhaps can our notions of when a trait is "relevant" to a particular

3. The Private Affirmative Action Issue

Notwithstanding the previous section's conclusions on disparate impact, the law currently permits an employer to adopt an affirmative action plan to ameliorate disparate impact and integrate her workplace. The *Weber* case¹⁰⁹ that so decided was surely controversial, and the majority's interpretation of the 1964 Civil Rights Act rather dubious in my opinion, but the result is clear enough. A bona fide voluntary affirmative action plan is not illegal discrimination against those dispreferred. If integration and equality of group status, rather than bias and aversion, are the employer's reasons for discriminating, the laws do not prohibit it.¹¹⁰

The moral case for voluntary affirmative action is not as certain as its legality. Voluntary affirmative action is rarely the product of bias. Working in favor of voluntary affirmative action is its potential to counteract the negative social effects of disparate group impact noted in the previous section. Factors working against voluntary affirmative action plans, however, include various negative social effects it may itself engender, particularly when sponsored by major private institutions. Ultimately, the moral case for or against voluntary affirmative action plans is an empirical one that cannot be settled from the armchair. Nonetheless, some potential negative effects of affirmative action that bear on its morality are worth mentioning. First, when major private institutions award jobs based on race, they tend to produce or exacerbate racial balkanization and racial politics. Racial groups will fight over their allocations of jobs. Some subgroups will demand their own allocation, while others will wish to be included within larger groupings.

Second, individuals will attend less to improving their own productivity and more to racial politics if racial group membership is an alternative to productivity as a means of advancement. Third, more productive groups will feel great resentment at being discriminated against because of their race. Civil strife born of resentment will replace the civil strife born of unequal group status. Fourth, negative biases toward and negative stereotypes regarding the beneficiaries will as likely be reinforced or increased in number as eliminated.

discriminatory choice.

¹⁰⁹ *United Steelworkers v. Weber*, 443 U.S. 193, 208 (1979).

¹¹⁰ See *supra* note 93 and accompanying text.

Fifth, the beneficiaries themselves may feel stigmatized, suffer self-doubt, and in order to protect their self-esteem adopt what Chris Wonnell calls the "affirmative action ideology," the claim that meritocratic values are racist or sexist.¹¹¹ Indeed, affirmative action beneficiaries will be more likely to adopt this ideology when the values behind the ordinary preferences that affirmative action trumps are deeply held by the affirmative action beneficiaries and others.

Finally, affirmative action raises the costs to the employer and thus to the consumer of satisfying ordinary preferences. Productivity is reduced in the sense that priorities are shifted from producing ordinary goods and services to producing more equality of groups.

Thus, private affirmative action is potentially a highly toxic form of remedy for the negative effects of disparate impact. Its own negative effects may be worse than those of group inequality. It is not intrinsically immoral, and it is currently legal, but it may be extrinsically immoral.

CONCLUSION

What makes wrongful discrimination wrong? I promised the answer would be messy, and it is. Discrimination may be intrinsically wrong because it is based upon biases, the incorrect judgments of lesser moral worth, or upon the shallow aversions or inaccurate negative stereotypes that are produced by such judgments. Discrimination may be intrinsically wrong because it is based on an unjustifiable ideology of moral role. Discrimination based on deep-seated aversions, accurate stereotypes, or reactions may be wrong, but it is not intrinsically so. Rather, particular types of such discrimination will be wrong in particular cultures, historical eras, and contexts, and not wrong in others. Discrimination resulting from preferences for goods and services, while wrong on occasion, is rarely so. Despite our history of abysmal race relations and the current state of racial strife as well as the long dominance in our culture of an ideology of gender roles, disparate impact along racial or gender lines rarely, if ever, warrants the conclusion that the ordinary preference-based discrimination that produces it is wrongful. The remedies for disparate impact, on the other hand—comparable worth and affirmative action—are on morally less sure footing.

¹¹¹ Wonnell, *supra* note 11, at 119-41.

In short, in an otherwise just society, discriminatory preferences are intrinsically morally wrong if premised on error, moral or factual, about the dispreferred. Discriminatory preferences are extrinsically morally wrong if their social costs are large relative to the costs of eliminating or frustrating them. And if a discriminatory preference is morally wrong—and if there is no moral right that protects its exercise—then there is a case for legally prohibiting its exercise if the costs of legal prohibition and enforcement are low relative to the social gains to be achieved.

Thus, I end with these whimpers and no bang. What makes discrimination wrong is usually quite complex as well as culturally and historically variable. If I have only sorted out the complexity and considerations accurately, I will be satisfied, for much discussion of discrimination fails to attain analytical clarity and methodological rigor. The topic always generates a fair amount of heat. It desperately needs more light.