

Discrimination, Disparate Impact, and Theories of Justice
Richard Arneson **word count 11241**

*{Final version of this paper was published in Sophia Moreau & Deborah Hellman, eds., **Philosophical Foundations of Discrimination Law**, Oxford University Press, 2013.}*

Contemporary ordinary common-sense morality strongly endorses the ideas that all members of a society ought to enjoy equality of opportunity and that equality of opportunity requires that the state should not discriminate against anyone on the basis of race or religion or skin color and that private individuals acting in a public capacity (roughly in the market economy) should also not engage in such discrimination.¹ These are thought to be basic moral requirements, so that, for example, engaging in racial discrimination is regarded as wrong in itself, not merely wrong because it brings about some further bad or fails to contribute to some further desirable goal. These are also thought to be imperative moral requirements, which ought to be enforced by law.

This society-wide moral consensus is, to speak, only skin-deep. When one seeks to clarify just what is being affirmed, disagreement and confusion appear.

The common morality and law concerning discrimination pose a challenge for act consequentialist moral theories. Act consequentialism says that what constitutes an act's being morally wrong is always that it brings about an outcome that is less good than the best outcome that might instead have been brought about. The challenge is sharper for welfarist consequentialist views, which hold that the value of outcomes is fixed solely by the aggregate well-being therein and its distribution across persons.² Regarding discrimination, common-sense deontology holds to the contrary that discriminatory acts that fit a certain description are morally wrong per se regardless of their consequences (at least up to some threshold level of bad consequences). My sense is that what I am calling the common morality here is deeply entrenched in enlightened public opinion in contemporary democracies.

In this essay I argue against this enlightened public opinion. I defend an act consequentialist approach to the law and morality of nondiscrimination. Part of the defense is to emphasize R. M. Hare's idea that moral thinking should be conceived as functioning on different "levels."³ Beyond some broad defense of the general family of act consequentialist principles, I single out a narrower family of views—prioritarian welfarist with welfare or individual well-being understood according to an Objective List construal—and suggest that the implications of the family for public policy in this area are especially attractive.⁴ Another element in my defense project involves exploring various suggestions as to what exactly the basic nonconsequentialist principles concerning discrimination might be and probing their inadequacies.⁵ This exploration focuses on the differences between disparate treatment, disparate impact, and accommodation requirements in the law and on their possible common rationales. To simplify discussion I use mainly examples from U.S. employment law.

A popular antidiscrimination norm holds that it is morally wrong to vary one's treatment of people on the basis of certain prohibited traits including race, skin color, religion, or sex. These traits should not play a role in the determination of what one does. This cannot be right as stated, because it would rule out passing sun screen to a light-skinned person rather than to a dark-skinned person. Suppose we restate the norm: it is

morally wrong to grant or withhold a benefit to anyone (or impose or decline to impose a loss) on the basis of the certain prohibited traits. This norm is advanced as a deontological rule, identifying a type of action engaging in which is wrong per se, independently of its further effects. When race and skin color are at issue, the rule prescribes that one act as though one were color-blind and unaware of racial distinctions. So call this the “color-blind norm.” So interpreted, an affirmative action or reverse discrimination program that requires one to favor members of groups that have suffered a history of mistreatment on the basis of race and are currently underrepresented in good jobs and slots of students in selective universities and other desirable social positions is straightforwardly in violation of the antidiscrimination norm and hence morally wrong.

However, many of us simply do not have the response that an affirmative action policy, just in virtue of offending against the colorblind conception of discrimination, thereby qualifies as morally wrong (or even pro tanto morally wrong).⁶ Affirmative action has uncertain, complex, and mixed effects, some good and some bad. Its evaluation is tricky. But it surely is not per se wrong in the way that rape, lying, breaking promises, and killing innocent nonthreatening persons who don’t want to be killed are thought to be per se morally wrong.

I shall suppose that one adequacy test for an account of nondiscrimination norms is that it should accommodate the thought that affirmative action policies are not per se wrong and can be a benign form of discrimination that is permissible and even required under some circumstances. Another adequacy test is that the account should explain why certain familiar bases of classification such as race, creed, and color are singled out for protection. The account should enable us to decide whether and to what extent the protection of nondiscrimination norms should be extended to further classifications such as sex, sexual orientation, age, disability, physical appearance, and physical attractiveness.

The color-blind antidiscrimination norm differs in its implications from a meritocratic norm: when choosing individuals to occupy desirable social positions, one ought to choose on the basis of the merits of the applicants for the positions. The meritocratic norm condemns choosing among job applicants capriciously or whimsically; the antidiscrimination norm under review does not. A stronger version of the meritocratic norm is *careers open to talents* (also known as *formal equality*).⁷ This holds that certain opportunities such as student positions in colleges and universities, desirable employment posts, and funds loaned by banks should be distributed in competitions open to all who wish to apply, with selection among applicants made on the basis of the merits of the applications, and merit being the degree to which awarding the opportunity to one versus another applicant would be reasonably expected to advance the morally legitimate purposes of the enterprise dispensing the opportunity.

All three of these norms fail the suggested adequacy test of allowing affirmative action policies to be permissible. One possibility here is that the norms do not state principles that hold without exception, but rather relevant moral considerations that might be offset by competing moral considerations. However, if we can identify exceptionless principles that match our intuitions concerning discrimination issues, this would surely be desirable.

Notice that as stated the colorblind norm condemns more narrowly targeted forms of discrimination: An employer who is perfectly open to hiring women and African-

Americans may downgrade the applications of aggressive women or dark-skinned African-Americans (but not the applications of aggressive men or dark-skinned Italian-Americans), and this practice should qualify as violating of the norm under review on the ground that acting so involves a disfavoring of a qualified sort on the basis of membership in the broad category of sex or race.

Animus or prejudice.

In this section and the next I search unsuccessfully for a deontological principle of wrongful discrimination. This principle would distinguish wrongful from nonwrongful acts of discrimination and justify the distinction.⁸ One proposal is that intrinsically wrongful discrimination occurs just when an agent treats a person identified as being of a certain type differently than she otherwise would have done because of unwarranted animus or prejudice against persons of that type.⁹ This proposal accommodates the idea that there can be wrongful discrimination when the agent is not consciously aware of the mental processes that constitute her making discriminatory choices. One might be moved by unconscious bias, a cognitive distortion that influences one's treatment of people. One might yet be at fault for having this bias and acting from it. The proposal also has no trouble accommodating the claim that there can be benign discrimination among people on the basis of their protected group membership that is nonetheless not morally wrong. Following the dictates of an affirmative action or reverse discrimination program, I can favor black applicants over equally qualified white applicants without acting from any sort of animus or prejudice against anybody on the basis of their belonging to one or another social group.

The proposal under review runs into trouble with examples of opportunistic uses of group membership distinctions for personal gain that do not proceed from any animus or prejudice but that nonetheless strike us as clear instances of wrongful discrimination. Suppose blacks are interested in moving into my residential neighborhood, and fearing that property values in this area may decline in consequence, I organize a committee dedicated to maintaining racial purity in the neighborhood and discouraging blacks from purchasing homes here. My motives may be simply profit-maximizing, and involve no racial animus or prejudice, but what I am doing still seems clearly wrong and wrongful discrimination. Or suppose a social norm against hiring Hispanics for skilled jobs depresses the job prospects of skilled Hispanics, so I can attract Hispanics to accept skilled jobs in my firm for lesser pay than I would offer to identically qualified whites applying for the same jobs. Again, I am simply exploiting a situation and driving a hard bargain, and my actions need not proceed from any sort of animus or prejudice against Hispanics or anybody else. I would be happy to do the same to skilled whites if the tables were turned. Nonetheless what I am doing looks to be wrongful discrimination.¹⁰

Demeaning and subordinating.

Deborah Hellman identifies wrongful discrimination as follows. Discrimination is classifying people on the basis of trait possession and treating them differently on that basis. Doing this is wrongful just in case it is demeaning. In turn *demeaning action* is understood as action that (1) expresses the view that a person one's action affects has less basic moral worth than others (lacks full equal humanity) and (2) is done by someone with power or status so that what is done constitutes putting down or subordinating the person.¹¹ Hellman distinguishes between acts that are wrongful qua discriminatory and acts that are discriminatory and also wrong but for reasons unrelated to their being acts of

discrimination (in other words, acts can be morally wrong and acts of discrimination but not instances of wrongful discrimination).

This characterization of wrongful discrimination accommodates the idea that there can be benign as well as malign discrimination and there should be no presumption that one engaging in the former acts wrongly. An affirmative action program that favors members of underrepresented groups in admissions to places in selective colleges may well not convey anything resembling a message that those candidates for admission that the policy disfavors are of lesser basic moral worth than others. The characterization under review also distinguishes between capricious, idiosyncratic discrimination, as for example refusing to hire job applicants with thick earlobes, and discrimination that targets traits such as skin color that are associated with a history of stigma and mistreatment and currently are a marker of low social status. The former type of discrimination is again unlikely to be expressing the view that those with thick earlobes have lesser worth.

What might at first seem minor problems of formulation mar the proposal. Suppose a society of devout Christians holds firmly that all people are equally loved by God and of fundamental equal worth and destined for an equally happy afterlife and should obey divine commands prescribing racial caste hierarchy during our temporary sojourn on earth. Surely acts by these Christians enforcing Jim Crow type segregation should qualify as wrongful discrimination even though these acts could not be claimed to express the view that those targeted for adverse treatment are of lesser basic moral worth than others.¹² Or suppose there is a society composed of black people and white people. The whites have a raw animus against the blacks. The whites do not claim that the blacks possess traits that merit negative appraisal; they simply react with repulsion to the sheer typical appearance of black persons. Nor do the whites believe the blacks have lesser fundamental moral worth. The whites simply hate the blacks and tend to act in ways that advantage whites over blacks. Again, in the scenario just described, we should judge the animus-based discriminatory behavior of whites toward blacks to be paradigm instances of wrongful discrimination, but Hellman's formulation does not allow this verdict.

It won't do to amend the proposal so that it would count as wrongful any action that expresses the idea that those one act affects are lesser in nonbasic worth than other people and satisfies the other conditions. This won't do because however exactly one construes the idea of nonbasic worth, some people will have less of it and some more, and actions that express the belief that this is so are not thereby rendered morally wrong. Some stereotypes may be accurate.

The doubts in the previous paragraphs concern condition (1). Condition (2) is also suspect. A powerless, low-status person who has the opportunity (as he thinks) to provide lifesaving aid to some accident victims, and deliberately refrains from extending any aid to members of social groups he hates and vilifies, just on that basis, even when extending aid would (as he thinks) be costless to other accident victims, is wrongfully discriminating, we should say, and the fact that as it happens he lacks the power effectively to channel aid to anyone does not affect the appropriateness of this moral judgment. At least, we do not want our notion of wrongful discrimination to rule out that pending further description, the discriminating behavior described can qualify as wrongful discrimination.

Perhaps these problems of formulation can be fixed. The fixing involves stipulating what it is to be demeaning that does better at fitting our considered judgments about what sort of discrimination should count as wrongful. Roughly, suppose we say that wrongful discrimination is discrimination that expresses a view about the targeted group that opposes the ideal of a society of democratic equality in which all people relate as equals, caste hierarchy is abjured, and no one's interests are discounted in the determination of appropriate public policy and individual action on a basis of animus or prejudice. Another possibility along this same line: wrongful discrimination is discrimination that expresses a view about the targeted group as just described and that is of a type that tends to hinder the emergence of a society of democratic equality. The former formulation is consistent with Hellman's insistence that in order to qualify as wrongful, discrimination need not harm anyone or have any negative effect beyond putting someone down, which does not entail that the person put down feels bad about herself or loses self-respect or a sense of her own basic worth or suffers any further loss of social standing or any other harm.¹³

These suggested norms fall under the general heading of rational attitude accounts of right and wrong action.¹⁴ Right actions express rational, appropriate attitudes toward those who are or might be affected; wrong actions express irrational, inappropriate attitudes. One general worry about such accounts is that an action may express a perfectly reasonable attitude yet harm, or fail unreasonably to help, some of those affected. My love for Fred may be perfectly reasonable but expressing it in action may just cause distress to no good purpose. Suppose we say it is a necessary, not a sufficient condition of being right, that an action must express rational, appropriate attitudes toward those who are or might be affected. This proposal falls to the ground when we consider that in some circumstances expressing an irrational or inappropriate attitude may be the only available way to produce significant good or avert harm. Insulting me by expressing an inappropriate attitude toward me may be the only possible way of inducing me to fulfill an important duty toward others.

It might be thought that case by case criticism of deontological norms is otiose once one has adopted a consequentialist standpoint, because this doctrine rejects all such norms once and for all. Hence any adherent of a thoroughly consequentialist and welfarist morality will have to reject any account of wrongful discrimination along the lines being considered. The welfarist consequentialist after all holds that acts are to be assessed as right and wrong according to some function of their impact on individual human well-being. If you describe a type of action in a way that leaves it open whether the action so described harms or benefits anyone, the description has to be morally neutral according to the welfarist consequentialist. Being an act of that type cannot render the act *pro tanto* either morally right or morally wrong.

Prioritarian act consequentialism.

The welfarist consequentialist, so it seems will have to dig in her heels and resist other accounts of wrongful discrimination besides the one currently under review. Along this line Kasper Lippert-Rasmussen dismisses the family of ideas that identify wrongful discrimination with discrimination that either proceeds from animus or from false factual or evaluative beliefs about the group that is being targeted for the short end of the stick of discriminatory treatment.¹⁵

As it happens, Lippert-Rasmussen espouses exactly the version of welfarist consequentialism I also find to be most promising—desert-catering prioritarian consequentialism. This is a version of act consequentialism that holds that one morally ought always to choose an act whose overall consequences are no worse than the consequences of anything one might instead have done. The measure of the goodness of consequences is total weighted well-being (that accrues to persons and other sentient beings—I leave aside the issue of how to balance the interests of persons and other types of sentient beings). A benefit one obtains for a person is better, the larger the well-being gain it brings about, and better, the worse off the person otherwise would have been in lifetime well-being, and better, the more deserving the person. This description characterizes a broad family of views; the best member of the family is the one that assigns the most appropriate weights to the elements of weighted well-being. Whether an act is an instance of theft, deception, killing of the innocent, discrimination against women, racial minorities, or the aged, and so on, matters morally to its being morally right or wrong only insofar as these characteristics cause weighted well-being to rise or fall. So, for the act consequentialist, discrimination can only be a hindrance or means to what matters in itself, and the question about what constitutes wrongful discrimination looks to be wrongly posed.

Levels of norms.

There is an important qualification to be noted here. Following various theorists of act consequentialism, in particular R. M. Hare and Peter Railton,¹⁶ I note that an adequate moral theory needs to distinguish distinct levels of moral thinking. The act consequentialist principle is a criterion of moral right and wrong or theoretical determiner of what features of an act constitute its being right or wrong. This leaves it open to what degree this principle ought to serve as a practical guide to decision making by the individual agent, and to what extent institutions and practices should be established for this purpose. In view of the fact that human beings tend to be selfish (to prefer favoring themselves and those near and dear to them rather than anyone who might be affected at any time or place by what they do), not well informed about empirical and evaluative facts that are material to forming correct judgments about what ought to be done, and not very competent at integrating such information as they do possess into the determination of what ought to be done, act consequentialism is usually a poor practical decision making guide. Being selfish, not well informed, and not good at reasoning, in many situations I will do better at choosing acts that conform to act consequentialist standards if I eschew direct calculation of what act of those I might do would lead to the morally best outcome and instead follow simple rules. What rules? There are different sorts of rules corresponding to different levels of moral thinking. From various standpoints it is commonly accepted that the legal rules enforced by the state should not perfectly mirror fundamental moral principles. Laws have to be coarse-grained, simple enough that those subject to law can figure out their requirements and generally designed to be implementable at reasonable cost. Even in a society whose legal rules were selected by act consequentialist calculation, the legal rules that ought to be in place would not be the single rule: do whatever would bring about the best outcome. For much the same reasons, the informal social norms that also regulate people's behavior ought to diverge from act consequentialism in their content.

Act consequentialist reasoning leads to the conclusion that laws and social norms should be established and that in deciding what to do one ought generally to follow laws and social norms rather than attempt to follow act consequentialist principles. Act consequentialist reasoning in just the same way dictates that there should be a public morality, a set of moral rules designed to guide people's decisions about what to do. To function properly, this morality needs to be accepted by members of society; we should be trained to accept the moral rules and employ them in regulating our own and other people's conduct. This would all be true in a society whose public morality rules were set by correct act consequentialist calculation. In any ongoing well-functioning modern society, there will be a public morality, which might be good, bad, or ugly by act consequentialist standards. Unless the society is completely off the rails and set on evil, the individual will generally do better, in deciding what to do, by following the given rules rather than trying to calculate what to do by applying the fundamental act consequentialist principle to one's particular circumstances. Deciding whether or not to be unfaithful to my wife, I am generally likely to act better (by act consequentialist standards) if I simply follow the relevant moral rule "Don't be unfaithful" rather than ask myself what would be best on the whole.

At a given level of moral thinking, the rules in play are better or worse by act consequentialist standards, if with these rules in place the shortfall between the consequences of what one actually does and the best consequences that one might have brought about, aggregated across all decisions affected by the rules, is smaller rather than greater. The smaller the shortfall, the better the rules.

One might question the coherence of act consequentialist morality as just described. The public morality rules that one has internalized and is using to guide one's decisions will likely often conflict to some extent with what a correct application of the act consequentialist principle to the decision at hand would specify. The public morality and the act consequentialist morality seem to be unavoidably in conflict, so one will get contradictions: I ought to keep my promise (according to public morality) here and now and I ought to break my promise (according to act consequentialist morality) here and now.

There is no contradiction in asserting both of these claims. Both could be true. But one might still wonder how a person might internalize or accept public morality and yet concede its dictates do not determine what is morally right and wrong.

A partial response is that accepting a morality includes becoming disposed to follow its dictates and to have emotional and judgmental responses that accord with the accepted views. Accepting the component of public morality that says deception is wrong, I am disposed not to lie, to regard instances of lying by others and by myself with moral disapproval, to feel bad about lying, to be disposed to react negatively in my behavior to those I suspect of lying, to judge that those who lie are behaving wrongly, and so on. All of that is compatible with also believing, as a matter of theoretical morality, that really I ought morally always to do whatever would bring about the best outcome. In certain situations, perhaps many, perhaps all, this theoretical moral belief does not impinge on my reactions and deliberations.

How should one deliberate and decide on what to do in circumstances in which one senses that that following the public morality rules that apply to this situation would conflict with conformity to act consequentialist principle? As Hare observes, that

depends on what sort of deliberator and agent one is, generally speaking and in this particular sort of decision problem. After all, what one senses to be true might yet be false, and even if one's hunch here is correct, it might still be that one is likely to do better by act consequentialist standards if one ignores act consequentialism here and now. Just as there is an act consequentialist answer to the question, should one now try to dispose oneself, or train others, to employ public morality rules rather than act consequentialist principle in deliberations about what to do, there is an act consequentialist answer to the question, should one now try to dispose oneself, or train others, to become and remain theoretically convinced that act consequentialism is the supreme moral standard and to revert back to using act consequentialism as a decision guide if cues in one's circumstances signal a sufficiently great disparity between what public morality tells one to do here and what act consequentialism would dictate. The answers might vary by degree from agent to agent and for different types of agents likely to face different types of decision problems.

Perhaps for some societies, and for some agents, the answer is blanket suppression of act consequentialism as any sort of guide to decision making and selection of actions. For these societies and for these agents, act consequentialism by its own lights ought to become self-effacing. If act consequentialism is true, this thought should be suppressed, on act consequentialist grounds. But consider the mental state of someone who (let us suppose, rightly, by act consequentialist standards) combines some degree of acceptance of the going public morality of one's society (or some variant of it of one's own devising) and the theoretical belief that the fundamental moral standard, the criterion of right and wrong, is act consequentialism, along with some tendency for this theoretical belief to intrude in certain types of situations on one's practical deliberations. This person's set of moral beliefs is likely to be a jumble: accepting public morality as morality, she believes that lying, breaking one's promises, killing innocent nonthreatening people who have a life worth living, and so on are morally wrong, and also believes that in any situation nothing is morally right except what would bring about best consequences (and that in possible and very likely actual circumstances, hallowed public morality rules will dictate conduct that would not lead to the best consequences). There is a possible fully consistent position that combines dispositional allegiance to public morality with full awareness that all genuine moral normativity flows from the fundamental level of act consequentialist principle. But adhering to this fully consistent position might be worse, from an act consequentialist standpoint, than adhering to some jumbled position. The act consequentialist theorist will say this is a problem of life, not a defect in act consequentialist moral theory.

How the consequentialist might embrace nondiscrimination.

The point of rehearsing these features of act consequentialist moral theory is neither to defend nor attack act consequentialism. I simply want to call attention to two possible stances an act consequentialist might adopt toward ideas of wrongful discrimination, equality of opportunity, the ideal of a society free of caste hierarchy, and so on. Suppose the act consequentialist says she rejects these norms. That might mean one of two very different things. It might mean merely, what has to be true, that act consequentialism fills up the space of fundamental moral norms and leaves no room at that level of moral thinking for deontological and other nonconsequentialist notions. Saying just *that* is fully compatible with upholding antidiscrimination and equal

opportunity norms as important components of public morality as it ought to be. The alternative possible stance of the act consequentialist toward a proposed antidiscrimination norm involves a more robust and thoroughgoing rejection. She might reject these norms not only as rivals to consequentialism but also as components of the public morality that is either defensible in present circumstances or defensible as part of the public morality for modern societies that would be best by act consequentialist standards. The more thoroughgoing act consequentialist rejectionist would hold something like the following view: I endorse a public morality that includes certain agent-relative constraints and agent-relative permissions, a public morality that condemns lying and promise breaking and theft and tolerates not giving all of one's wealth and income to Oxfam, but this public morality endorseable by act consequentialist morality does not include antidiscrimination and equal opportunity norms.

If the welfarist act consequentialist position is on the right track, and people are to some extent implicitly following it in grappling with issues about discrimination and equality of opportunity, it should not be a surprise that that these issues continue to be puzzling and difficult to resolve even for people of good will who share broadly liberal sympathies. If the act consequentialist is right, what shape the social norms and law and public morality of nondiscrimination ought to have depends on difficult evaluative and empirical questions concerning what set of policies would maximize priority-weighted well-being over the long run. To put it bluntly, we do not know the answers to these questions, so we are not in a position firmly to identify the morally best set of nondiscrimination norms and practices. Even if common-sense morality is not implicitly quasi-utilitarian as Henry Sidgwick surmised, common-sense morality may be just going off the track insofar as it claims that staring ever more intently into the deontological pool of moral claims will yield the right answers (or has already done so).¹⁷

My view about what act consequentialism implies regarding discrimination crucially depends on the particular sort of act consequentialism that is being affirmed. Notice in particular that the priority version is welfarist and also (1) supposes that interpersonal comparisons of well-being make sense in principle and can be made at least sometimes and in a rough way in practice and (2) declines to identify welfare or well-being or the good life for a person with preference satisfaction but instead identifies the good for an individual with particular types of achievements (the items on a so-called Objective List).¹⁸ If one is a welfarist act consequentialist and identifies welfare with preference satisfaction and denies interpersonal comparison, one is likely to end up affirming only the Pareto norm (it's wrong to bring about or tolerate states of affairs which can be changed by making somebody better off in preference satisfaction without making anyone else worse off), then immediately any equal opportunity or nondiscrimination norm becomes either unstateable or normatively problematic.¹⁹ One view is that the state ought not to require individuals to engage in discrimination against currently protected groups but should allow individuals to do whatever they wish, including discriminate, with their own property and person. From this austere standpoint—nothing matters except preference satisfaction, which cannot be measured across persons-- on what basis should the state favor the satisfaction of the black who seeks opportunity for employment and public accommodation in restaurants and stores rather than the preference satisfaction of the people who want to exclude blacks on racist grounds? Richard McAdams, defends state prohibition of such discrimination from the

austere standpoint. His argument is that racial subordination arises from group status competition, and since status is relative, whenever one person rises in status, another must fall. Hence the deployment of resources in such status competition can be socially wasteful, and state regulation prohibiting discrimination can reduce this inefficiency and be justified on that ground.²⁰

In contrast, prioritarian consequentialism can support nondiscrimination practices and laws in circumstances in which the spare efficiency argument cannot be sustained. Priority can say that satisfying racist preferences does not per se enhance one's well-being and make one's life go better, and can add that when victims of discrimination tend to be among the worse off members of society, there is special moral urgency to bringing about genuine gains to their well-being. (A desert-catering version can add that being disposed hostilely toward those of disfavored race or sex or the like can render one morally undeserving and so less morally eligible for state action to boost one's well-being.)

Priority can possibly defend nondiscrimination and equal opportunity norms as part of the best consequentialist public morality. But will it, in our circumstances? I say Yes, but lack knockdown arguments. That welfarist act consequentialism fails unambiguously and certainly to specify the appropriate public morality and legal treatment of discrimination issues in the absence of lots more empirical knowledge than we will soon possess is not per se an objection to act consequentialism. Maybe our stance in this domain should indeed be tentative and uncertain. If the knowledge consequentialism implies that we need would resolve our perplexity if we could obtain it, that is a point in favor of this doctrine. Only if after reflection we find we are committed to affirming a morality of nondiscrimination for reasons that consequentialism does not register would we have here the makings of a case against the acceptability of consequentialism.

On these issues, the consequentialism here affirmed does not leave us entirely in the dark. It is plausible to affirm, as a component of public morality, the ideal of a society of democratic equality in which all people relate as equals, caste hierarchy is abjured, and everyone's comparable interests count equally in the determination of appropriate public policy and individual action.²¹ The plausibility of this ideal hinges on the plausible conjecture that compared to alternatives, democratic equality fosters good quality lives for people, with the good in lives fairly distributed across individuals. The question immediately arises, how much equality should democratic equality demand? As Samuel Scheffler notes, equality may be a plausible ideal but it is also puzzling, since significant inequalities and hierarchies are rife in modern democratic societies and we surely do not want to strike down all of them.²² Priority says here that we should insist on democratic equality in relationships just when (and to the degree that) doing so is productive of good lives fairly distributed.

Defending priority on discrimination.

By assigning only instrumental moral significance to the phenomenon of discrimination, priority runs against common reflective opinion. Following Hare, my response is that it may be, depending on the facts, that in some settings priority may itself urge that people establish laws and internalize norms that prohibit certain types of discrimination, up to a point, no questions asked, and in particular, independently of

expected consequences. So a partial accommodation of common opinion may lie along this path.

Beyond accommodation, there may be ways to undermine common reflective opinion. One undermining strategy is to show by example that equality of opportunity norms allied with nondiscrimination norms could be violated without the violation striking us as even pro tanto wrongful. Consider an imaginary primitive society that lacks administrative capacity and settles many matters concerning how to live by adhering to fixed conventions. In this society by customary rule men are assigned to the role of hunter and women to the role of gatherer. On the whole, gatherers live better than hunters (longer lives, better health, greater fulfillment). There is no procedure whereby conventional role assignment might be altered by a showing of individual competence; suppose that the costs of developing and instituting such a procedure and the social tensions that operating it would provoke would exceed the benefits as evaluated by the prioritarian standard. For similar reasons there is no cost-effective way to institute a redistribution scheme that would compensate men for the lesser benefits they gain, compared to women, from the fixed division of labor. The existing scheme, let us stipulate, is optimal according to the prioritarian standard. (Notice that the assumptions needed to drive this result are somewhat far-fetched.) According to priority, equalizing welfare across persons is not in itself morally valuable, nor is equalizing across social groups. Priority is indifferent to the massive violation of the norm of careers open to talents in the imagined society, and would be similarly indifferent if men were systematically advantaged compared to women.

The example also illustrates opposition between priority and the range of equal opportunity norms, not only careers open to talents but also more demanding principles including fair equality of opportunity (all those with the same native talent endowments and the same ambition to develop and exercise them should have the same prospects of success) and luck egalitarian equality of opportunity (it is morally bad—unjust and unfair—if some are worse off than others through no fault or choice of their own).²³ I hope that the example will elicit agreement that in the stipulated circumstances, when violations of equal opportunity and nondiscrimination norms help to make people's lives better according to the prioritarian standard, there is no moral loss. These violations are not pro tanto or even prima facie let alone all things considered morally wrong.

What classifications should the nondiscrimination principle single out for protected status?

Many of us hold that it is wrong to discriminate on the basis of a person's sexual orientation in just the same way that it is wrong to discriminate on the basis of race, creed, color, or sex. One may wonder about further extensions of the scope of protection against discrimination? Age? Disability? Physical appearance? Physical attractiveness? In January, 1992 the City Council of Santa Cruz, California considered a proposed city ordinance that would prohibit discrimination against individuals on the basis of personal appearance. In the context of employment, such an ordinance would rule out favoring or disfavoring an applicant for hiring or promotion on the ground that she is physically attractive or unattractive, and also on the ground that she conforms or fails to conform to conventional standards of dress or appearance. (These rules would be qualified by allowing that personal appearance be a factor influencing employment decisions when it is a bona-fide occupational qualification (BFOQ), for example, when a certain

appearance is required in order to carry out essential functions of the job for which one is applying).²⁴

Whether the law should single out a classification for protected treatment depends on many considerations, including considerations of administrative practicality. But we might wonder whether there is a background principled basis for holding that discrimination against people on the basis of their possession of certain traits or membership in certain groups is specially, intrinsically morally wrong. We might say that arbitrary or unjustified discrimination is wrong, but this is uninformative and unhelpful. We might also wonder whether there is a background principle that tells us how to determine whether a candidate classification deserves inclusion in antidiscrimination law. In this connection physical appearance indicates the problems. Some examples of discrimination on the basis of appearance that we can imagine appear to be paradigm instances of wrongful discrimination and other examples not so and how to draw appropriate lines here is not, to me anyway, at all obvious.

Different classifications such as age and sexual orientation raise different concerns, and we might be skeptical that there is a one-size-fits-all moral principle that encapsulates the grounds for distinguishing acceptable and unacceptable discrimination. Sophia Moreau suggests that the grounds for including a classification under antidiscrimination norm protection may indeed be disparate, and that what unites the category is rather that we should bring it about that the possession of certain traits (religion, race, ethnicity, sex, sexual orientation, and so on) by an individual should be factors she does not have to regard as imposing costs when she is considering participation in public sector activities such as applying for a job, deciding what restaurant or store to patronize, and so on.²⁵ Any such costs are costs that people in general participating in the activity will absorb as an accommodation to the trait possessors. Nondiscrimination law and social norms carve out deliberative freedoms as just described for possessors of traits deemed deserving of special protection on various grounds. The trouble with this proposal is that I do not see that such across the board accommodation is ever a good idea. Choosing a religion, for example, involves acceptance of myriad requirements and permissions that may affect in an indefinite number of ways the benefits and costs of seeking one or another form of employment. There is no basis in general for thinking individuals should be insulated from such costs or should have the deliberative freedom to choose their religion without any consideration of the costs their choice might impose on others in various circumstances and how those costs might fairly be spread among persons or confined to the individual cost-generator. With respect to unchosen traits, we should expect that costs should often fall on the person who can most easily or cheaply minimize them, which may often be the trait-possessor not others.

Here I shall simply reiterate a suspicion already voiced. I do not see a principled basis for deciding what types of traits and group classifications nondiscrimination practices should be concerned to protect except by looking to the consequences of extending and denying protection to candidate classifications.²⁶ Pondering what types of conduct are per se intrinsically wrong is not fruitful. This point takes us a step toward acceptance of the consequentialist approach but does not commit us to that. One might hold that no type of act qualifies as wrong unless its description entails that it does harm, reduces someone's well-being (or fails to increase well-being when boosting is morally

required), but deny that all things considered calculations of consequences determine where to draw the line between the permissible and the wrong. (I do not see much future for this intermediate proposition, but nothing in this essay rules it out.)

One might also draw a positive lesson for a deontological account of the rights and wrongs of discrimination from this discussion. Morally objectionable discrimination is a diverse phenomenon. There is unlikely to be one deontological principle that holds always and everywhere and states necessary and sufficient conditions for wrongful discrimination. A more promising alternative is that there are wrong-making characteristics of discrimination, such that if an act of discrimination embodies any of these characteristics, its doing so is a pro tanto consideration against its moral permissibility. These characteristics might include being demeaning as specified by Hellman, expressing an attitude expressive of caste hierarchy, being done from animus or prejudice, and so on. These characteristics can be outweighed by countervailing factors, and whether a given act of discrimination is wrong all things considered depends on the overall balance of considerations. An account along this line might be correct; the doubts and objections I have raised are not decisive against it. We might regard such an account as a fallback position, to which we might have to retreat if efforts to arrive at a more systematic principled position fail. Priority looks to be a horse that is very much in the running in this competition among candidate systematic principle accounts.

The U.S. employment discrimination law.

In U.S. law (as in that of other countries), antidiscrimination provisions are diverse. The Establishment and Free Exercise clauses of the First Amendment to the Constitution prohibit government from discriminating among citizens on the basis of their religious beliefs or affiliation. The Equal Protection clause of the Fourteenth Amendment prohibits government from acting in a way that denies equal protection of the laws to any citizen on the basis of race or creed or color or national origin.

Federal laws prohibit discrimination on the basis of age, and in particular discrimination that benefits younger people at the expense of older people, in certain settings. Another federal law prohibits discrimination against individuals with disabilities and requires employers to provide reasonable accommodation to prospective and current employees who have disabilities provided this can be done without undue hardship to the enterprise. An example of such an accommodation would be providing a translator fluent in American Sign Language to assist a deaf professional in conversing with clients and associates.

To simplify discussion, I focus on antidiscrimination and equal opportunity norms in employment law, and specifically on U.S. federal law. Even on this narrow terrain, my description is incomplete and stylized. I restrict attention to U.S. law in this area not because it is exemplary or emblematic, but simply because I lack the competence to make comparisons across laws in different countries.

U.S. employment law forbids disparate treatment of protected groups in employer decisions about hiring, promotion, and conditions and benefits of work. Disparate treatment involves, for example, denying an applicant a position for which she is qualified because of her race.²⁷ The law also regulates disparate impact.

The disparate impact component of the law works as follows. If an employer uses a hiring procedure that has a disparate impact on individuals who are members of protected groups, defined by race, color, religion, sex and national origin, that establishes

a prima facie case. Disparate impact here means that the proportion of protected-category applicants who are hired is smaller than their proportion of the relevant labor pool. If the employer is sued, and a prima facie case is established, she can rebut the prima facie case by showing that the hiring procedure in question is job-related and justified by business necessity, unless the government agency or individuals bringing suit can propose an alternative hiring test that is just as good for the purpose and would not have such disparate impact.²⁸

The justification of disparate impact law.

Legal theorists and philosophers have disagreed on the question, do the disparate treatment and disparate impact parts of antidiscrimination law rest on common moral foundations or are they morally discontinuous? A related question is whether the accommodation requirements in laws prohibiting discrimination against the disabled presuppose the correctness of stronger and more controversial moral principles than those to which one must appeal to make the best case for the rest of the nondiscrimination legal code.

The answers depend both on what is the best interpretation of what these components of the law are doing and on what are the correct moral principles that apply to these domains of law and determine their proper content. Not easy questions.

Suppose that a city uses the scores on a written exam as the basis for hiring firefighters and for promoting firefighters within the ranks. An applicant must attain a threshold score to be considered further, and within the pool of applicants, the cutoff disproportionately eliminates African-American and Hispanic applicants and leaves white applicants still in the running. One possibility here is that the situation involves disparate treatment by indirect means. The city administrators either intend to favor white applicants over the others and select the test just in order to bring about this result, or they are cognitively biased against the minority applicants, and they select the test, thinking it accurately gauges fitness for employment or promotion in this job category, whereas in actuality the test results do not correspond to applicants' varying abilities that are relevant to job performance. If this is the case, it is intuitive to suppose that disparate treatment and disparate impact rules are close comrades engaged in a common struggle. No great normative gulf separates them.

However, the question then arises whether disparate impact rules so construed are otiose. A sensible law against disparate treatment would allow a case to go forward if a hiring practice results in disparate impact and there is no plausible explanation of the employer's behavior other than that she is declining to hire because the applicant belongs to a protected group. So, why disparate impact?

The rationale for disparate impact must then take the form of a rationale for a mild form of affirmative action, as follows. Among the employment practices one could follow that would be about equally good from the standpoint of business efficacy, one is legally required to adopt the one that has the least negative disparate impact on members of protected groups. Disparate impact law so construed then can serve any of the social goals that affirmative action might be thought to serve. Disparate impact only involves a mild form of affirmative action because as described its implementation need involve no violation of the norm of careers open to talents.

Conservative critics of disparate impact reject the idea that it is compatible with careers open to talents and other norms they embrace. To appreciate their worries, go

back to the example of hiring firefighters. Accused of perpetrating illegal disparate impact, it might be the case that the city administrators do not intend to favor white applicants over others and harbor no cognitive biases against minority applicants. They sincerely believe, and have some credible reason for believing, that the test they are employing that results in disparate impact is a fair test of job fitness. If some applicants eliminated from consideration for hiring or promotion by the test sue the city for disparate impact violation, the city loses its case on the facts as so far specified unless the city can show that the test in question reliably sorts applicants according to a qualification relevant to job performance and the plaintiff does not in reply propose an alternative selection procedure that would not have disparate impact and the defendant does not show that the proposed alternative would not adequately sort applicants by ability. If making these determinations were certain and costless, the enforcement of the disparate impact rule would never bring about the result that the defendant is required to reduce disproportionate impact in hiring and promotion by selecting a less qualified minority over a more qualified nonminority applicant. Disparate impact enforcement would never issue in violation of the norm of careers open to talents in order to bring about a proportionate racial draw from the pool of applicants.

But suppose that there are significant costs associated with the task of presenting legally convincing evidence that the hiring procedures one employs are better than alternatives at selecting the best candidates, and in particular better than alternative procedures that might be suggested that would lessen the disparate impact of the procedures actually being followed. Suppose that there is a gray area, and within the gray area the state agency or private firm cannot demonstrate, or cannot demonstrate at feasible legal cost, the superiority of its actual hiring procedures, which have disparate impact, and have triggered legal challenge. In such cases the enterprise might be buffaloeed by disparate impact enforcement to give up using its chosen (and we are assuming, superior) hiring procedures and to accept inferior procedures that reduce disparate impact or to institute a de facto quota system in hiring that guarantees slots to members of protected groups to eliminate disparate impact. In these cases disparate impact might still be serving some social goals but would not be reducing discrimination against protected group members and would be bringing about violation of the formal equality of opportunity (careers open to talents) norm.

Notice also that there can be patterns of actual discrimination that vigorous enforcement of disparate impact legal provisions might exacerbate. Consider the example of a job market in which one protected group is underrepresented and another racial group, let us say Asian-Americans, is overrepresented. We stipulate that in fact there is discrimination in this market against Asian-Americans. An Asian-American applicant with relevant characteristics that are identical to the characteristics of applicants from other social groups will fare less well in the particular job market. In this scenario, disparate impact law is triggered by the underrepresentation of underrepresented protected groups under current hiring procedures, and if enforcement disparate impact has any effect at all, it will be to induce the employers to adopt practices that reduce disparate impact. In this example, these new procedures are likely to have the unintended effect of increasing the discrimination, the steady violation of careers open to talents, that Asian-Americans already suffer.

An even more ethically problematic scenario might be unfolding in the just imagined circumstances. The effect of discriminating against Asian-Americans might be intentional, in the following sense: if those enforcing the law are intent above all on bringing about a world in which in each significant job category, the members of all currently underrepresented and historically disadvantaged groups are represented in full proportion to their numbers in the relevant segment of the labor market, come what may, then in the hypothetical circumstances specified the achievement of his goal must mean discrimination against Asian-Americans.

Of course, these are objections not to disparate impact laws in principle, but to hypothetical implementation that has gone awry from the standpoint of disparate treatment and disparate impact themselves. Perhaps to some degree such problems of implementation can be alleviated by shifting the burden of proof. We might propose that the law specify that the a designated federal agency of the state funded by nationwide taxation revenues should bear the burden of establishing whether an employment practice challenged on disparate impact grounds is effective in sorting qualified from unqualified applicants and whether alternative procedures are available that would do about as well at this sorting task with less disparate impact. Under this regime, no business owner would have to bear the expense of proving the efficacy of its current procedures; nor need a challenged government agency divert funds from its budget for this purpose. Critics will have a raised eyebrows skeptical response to this particular proposal and more generally will hold that in practice the effect of disparate impact laws will be to impose a variously strong affirmative action program that will be a brake on economic efficiency and will violate careers open to talents.

Adoption and implementation of disparate impact law that governs employment might then bring benefits and incur some costs. On balance, is disparate impact law defensible?

From some standpoints, disparate treatment and disparate impact laws are both wrong, and wrong for essentially the same reason. A Lockean libertarian holds that each person has the moral right to do whatever she chooses with whatever she legitimately owns, provided she does not thereby harm other in certain specified ways, by force, fraud, theft, violence, or threat of these. If she chooses to allow (for example) whites only to work at skilled jobs in her factory, she might well be morally criticizable, but since she has the right to do this, it would be morally wrong, a violation of her rights, to force her to desist from this behavior.²⁹

The preference satisfaction consequentialist who eschews interpersonal comparisons of welfare will have trouble embracing disparate impact even if she supports disparate treatment laws as state action to eliminate wasteful competition for racial and sexual and similar status hierarchy. Only if implementation of disparate impact imposes no more than what I have called “mild” affirmative action will the efficiency case for it be clear.

If one is strongly inclined to believe that the efficient operation of a competitive market relies heavily on the uncodifiable savvy of owners and managers of business firms, one will likely stress the damage to economic productivity (which makes all of us better off in the long run) that energetic enforcement of disparate impact threatens to cause. Challenged by disparate impact lawsuit or threat of that, businesses will abandon

sound hiring and promotion policies and substitute policies that mimic rigid quotas in their effects.³⁰

One who upholds an ideal in the democratic equality family on deontological grounds might well doubt the empirical claims just made but will insist that even in the worst case scenario in which these claims are true, disparate impact might be justified. If the democratic equality ideal is a component of social justice, and justice trumps economic efficiency, then the fact that establishing democratic equality might bring some losses in economic productivity does not in itself count as anything close to a decisive objection to disparate impact law regarded as a means to democratic equality or partly constitutive of it.

How likely is this worst-case scenario anyway? The question introduces contentious issues, on which people of different ideological leanings are likely to disagree in the absence of compelling empirical evidence that is hard to find. At one extreme, one might believe that the hostility and prejudice that disparate treatment laws target play little role in present-day economies in advanced democracies, so that if a given set of hiring practices in a firm or industry results in disproportionate hiring of some groups over others, that simply reflects the greater aptitude for employment of the members of favored groups. At the other extreme, one might believe strongly that prejudice against disfavored social groups (some racial minorities, women, gays and lesbians and transgender persons, the aged, the disabled, and so on) is rife, and further that membership in any of these groups is very unlikely to be correlated with possession of traits that employers of firms and customers of businesses legitimately want. On the latter view, disparate impact will reliably indicate the violation of the moral rights of job applicants to equal consideration, rights that ought to be legally protected. For example, if one-third of those in the applicant pool for New York City firefighters are African-Americans, but less than one-third of African-Americans are being hired, year after year, that fact reliably signals rights violations of the African-Americans.

In the long run these empirical issues should be resolvable by empirical inquiry. Addressing them is beyond the scope of this essay. I venture one comment. In modern multicultural democracies, different social groups sustain somewhat different ways of life for their members. In societies that welcome immigrants and do not insist on assimilation to mainstream culture, different immigrant groups may preserve ethnic and tribal features, which include unique patterns of customs, values, and norms, which influence the upbringing of their children. In this environment one should expect that different groups will vary in the socialization they impart to their members that has a bearing on members' fitness for various jobs and professions. One ethnic group may stress ambition, perseverance, and mastery of complex technical skills, another may stress religious faith and loyalty to traditional ways. These differences in culture and socialization may affect the likelihood that a given individual member of the group will seek a career, say, as medical doctor, scientist, or lawyer, and may affect the skill level for that career that the individual will develop if she does have the ambition to pursue that career. This being so, we should not be surprised if some groups are not represented in the labor pool for some desirable occupations in proportion to their numbers in the population. Nor should we be surprised if the members of some groups who are in the employment pool for some occupation are disproportionately well qualified for that occupation compared to members of other groups in the same pool.

These facts tend to provoke unease. Many of us would like to see a society in which the accidents of birth into one or another social group do not affect the likelihood that one will develop ambitious life goals and achieve them. The Rawlsian principle of Fair Equality of Opportunity (FEO) encapsulates part of this concern. FEO holds that positions of advantage in society that give people greater shares of resources (that Rawls calls “primary social goods”) than others have should be open to all such that all those individuals with the same ambition and the same native talent should have exactly the same prospects of success in competitions for these desirable positions.

In the multicultural society I have described, FEO is likely to be violated, so disparate impact results. If parents of Hmong immigrants and parents of Chinese immigrants impart differential skill and work attachment to their children, then Hmong youth and Chinese youth with equal ambition to attain desirable career posts will have different prospects of success. Moreover, we might also feel that it is unfair if some people through accidents of birth do not enjoy the same chances that others have of developing ambition for career success. After all, having such ambition is usually a necessary means to success; a person is very unlikely simply to fall into success. Failure to develop ambitious desires and life aims may be just one sad result of a deprived childhood.

If the characterization of multicultural society I have given is accurate at least in a stylized way, what is the right response? Here people looking at the matter from opposed moral standpoints will disagree. We might urge that the impact of multicultural differences on socialization must be overcome or offset, so that FEO, a requirement of justice, is fulfilled. Public schools and government programs to enhance the development of children can be oriented so they help those from disadvantaged social groups and counteract any adverse influences they receive from inept and ill-advised family members and neighbors. An alternative to offsetting family influence is to nip it in the bud, by interfering with the social group’s capability to socialize members’ children into its values and customs.

Rawls himself proposes FEO as a moral constraint on allowable inequality. According to Rawls, unless FEO (and other conditions Rawls advances) is met, an equal distribution of primary social goods across all members of society should be maintained, FEO could be upheld independently of careers open to talents, but Rawls treats the two together as one principle, and I shall do the same here. We could reconstrue FEO as a social goal to be achieved, a component of a society that avoids the moral evils associated with wrongful discrimination. (But notice that FEO could be fulfilled in a society even if there is wrongful discrimination across the board, just so long as the discriminatory acts balance out in their effects, so that all those with the same native talent and same ambition have the same prospects of competitive success. Also, absence of wrongful discriminatory acts in society is compatible with violation of FEO.)

Rawls himself does not fully spell out how to resolve the conflict between the liberty to participate in social groups (including the liberty of parents to raise their children as they see fit and order their family life as they choose). One might hold that the freedoms to participate in one or another distinct culture in a multicultural society take precedence over FEO. On this view we ought to implement FEO, if at all, only so far as this can be done without restricting multicultural liberties. On the other hand, other nondiscrimination norms and practices including careers open to talents set limits to

multicultural participation liberty: in certain public action contexts, one is not permitted to favor members of one's own tribe or group, but is required to treat all comers impartially.

Notice that one might also deny there need be conflict between multicultural participation liberty and FEO. One might observe that at least in principle the fullest multicultural liberty can be maintained and FEO can also be fulfilled. In this regime no one is forbidden to favor her own family members or members of cultural groups to which she is loyal over outsiders. Since individuals command different amounts and quality of resources, this favoritism will tend to give some individual advantages in life, a leg up in competition, that others lack. Society, all of us together, committed to FEO, takes steps that exactly offset these acts that specially favor some over others. Asian-Americans may work to instill special diligence and ambition and mathematical and analytical skill in their young members, and in response society sponsors Head Start to the Max programs that fully offset any advantages Asian-American children would be gaining, and so on across the board. This would be prohibitively costly no doubt, but could be done.

Insofar as multiculturalism and cultural diversity lead to a distribution of desirable employment and entrepreneurial activities that is unequal across social groups, we should not expect the enforcement of disparate treatment and disparate impact employment law to play an important role in resolving the problem. At least, this is so if disparate impact is used as a tool for eliminating what is really hidden disparate treatment. If on the other hand disparate treatment is used as a form of strong affirmative action, namely insistence on relaxation of meritocratic standards in order to increase the proportion of members of disadvantaged protected groups who enjoy desirable employment and entrepreneurial posts, disparate treatment can help resolve the problem if strong affirmative action can be of assistance. There are various familiar ways in which this might be hoped to work. One idea is that if a thumb is put on the scale that weighs applications in order to favor applications by members of underrepresented groups, their presence in the desirable positions will encourage ambition formation in the underrepresented community and help eventually bring about circumstances in which members of these groups are just as likely to want desirable posts and be qualified for them as members of already advantaged groups.

Priority, again.

The point of rehearsing details of the fair equality of opportunity ideal and its connection to multiculturalism is to reiterate a position to which this essay keeps returning. To what extent should FEO be upheld come what may? How should we balance multicultural liberty and equal opportunity goals? To these questions the prioritarian consequentialist has a simple, clear, to my mind attractive answer: The standard for resolving these questions is what mix of policies and priorities would do most to maximize good lives for people with good fairly distributed across people and fairness encapsulated by priority to the worse off.

How far should disparate impact laws press in their mild (or in some circumstances strong) affirmative action components? Again, priority suggests an answer that is somewhat outside of the box from the perspective of current discussions. In principle, priority could balk at disparate impact implementation necessary to

implement FEO, and in other circumstances could insist on implementation of disparate impact even in the teeth of FEO.

Notice that under some circumstances priority can justify the operation of disparate impact laws when that amounts to implementation of affirmative action or reverse discrimination policies that conservatives tend to abhor. If aggressive enforcement of disparate impact laws improves opportunities for members of disadvantaged groups, and especially if the effects are diffuse and tend to trickle down to worse off members of the group, the resultant well-being gains, weighted by priority, can exceed any losses that result from the lowering of meritocratic standards.

For that matter, under some circumstances priority can justify rigid quotas that assign desirable opportunities to members of protected groups in proportion to their numerical share of the population. For example, imagine that relations between French-speaking and Flemish-speaking Belgians have become rancorous, so that trust between the groups is eroded. A quota system that reserves desirable public sector jobs for each group in proportion to its population numbers, with meritocratic selection procedures applying only within each nationality group pool of applicants, may be better than any alternative allocation of these jobs, as assessed by the prioritarian consequentialist standard.

However, in such circumstances the prioritarian would be acting not from deep commitment to nondiscrimination and equal opportunity ideals but from strict indifference to them (at the fundamental level, not the level of public morality). The deontologist who holds that we should accept bad consequences if need be in order to respect deontological constraints, with nondiscriminating and equal opportunity norms included among the constraints, will still draw a line in the sand and disagree.³¹

¹ . Why only in a public capacity? If I choose not to befriend a person who is otherwise appropriate for friendship merely on the ground that I dislike her race or skin color or harbor a general prejudice against women, that is surely a paradigm instance of wrongful discrimination.

² . What the text refers to as “welfarist consequentialism” is what T. M. Scanlon calls “philosophical utilitarianism” in his “Utilitarianism and Contractualism,” in Amartya Sen and Bernard Williams, eds., *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982), 103-128.

³ . R. M. Hare, *Moral Thinking: Its Levels, Method, and Point* (Oxford: Oxford University Press, 1981), esp. chap. 2-3.

⁴ . On priority, see Derek Parfit, “Equality or Priority?”, reprinted in Matthew Clayton and Andrew Williams, eds., *The Ideal of Equality* (New York: Macmillan and St. Martin’s Press, 2000). On Objective List accounts of individual well-being, see Derek Parfit, “What Makes Someone’s Life Go Best?”, Appendix I in his *Reasons and Persons* (Oxford: Oxford University Press, 1984), 493-502.

⁵ . See Richard Arneson, “Against Rawlsian Equality of Opportunity,” *Philosophical Studies* 93 (1999), 77-112; also Arneson, “What Is Wrongful Discrimination?”, *San Diego Law Review* 43 (2006), 775-808.

⁶ . The term “affirmative action” is used broadly. Here I mean to refer to what is sometimes called “reverse discrimination.” Affirmative action in this sense occurs when in order to increase the representation of members of underrepresented protected groups in the set of those who gain some competitive good, meritocratic norms for selection among candidates for the good are set aside or overruled. Affirmative action might take the form of introducing quotas or putting a thumb on the scale in favor of candidates from underrepresented groups. Affirmative action policies can vary by degree, playing a trumping role in selection or having more or less weight as one consideration among several. (If one say gives more credit to a certain SAT score earned by a minority applicant than if it had been earned by a majority applicant because the score of the former involves overcoming special obstacles and indicates greater merit than the same score earned by the latter, that policy is a specification of meritocratic norms not an overturning or ignoring of them.)

⁷. *Careers open* as characterized in the text is stronger than the meritocratic norm just in being committed to a specific account of what makes an application meritorious. The account as stated is imperfect. Suppose that Sally is demonstrably better at accounting tasks than her competitor Joe, but hiring Joe rather than Sally for an accounting post will do more to advance the morally legitimate interests of the firm, because Joe has many wealthy connections and his presence in the firm will do more to attract lucrative clients than would Sally's presence. One might still hold that Sally is more qualified.

⁸. A complication is that some acts that do not count as acts of wrongful discrimination may nonetheless be wrong for other unrelated reasons.

⁹. See Arneson, "What Is Wrongful Discrimination?" for a qualified defense of the proposal, which derives from an essay by Larry Alexander, "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies," *University of Pennsylvania Law Review* 141 (1992). The proposal identifies wrongful discrimination with acting from unwarranted animus or prejudice, and the inclusion of the normative term renders the proposal incomplete, not very informative.

¹⁰. Another objection against the proposal is that it renders the idea of wrongful discrimination trivial, in that idiosyncratic disfavoring of people with a trait such as having large earlobes from animus or prejudice against members of this group would count as wrongful discrimination. I see that such discrimination would be unlikely to have large negative effects, unlike discrimination against people whose traits have spurred a history of oppression. But insofar as discrimination ever seems intrinsically wrongful on deontological grounds, the imagined hostility against those with large earlobes would appear a clear instance. Another objection is that the proposal conflates

having a bad motivation with acting wrongly. One can act with wrongful animus without doing anything wrong, as when a thug whose motivation is to do whatever it takes, even murder, in order to get a pack of cigarettes quickly, finds that the best means to his end is just to pay the posted amount for the pack of cigarettes to the convenience store clerk (Derek Parfit's example). I try to respond to this objection in "What Is Wrongful Discrimination?"

¹¹ . Deborah Hellman, *When Is Discrimination Wrong?* (Cambridge, MA: Harvard University Press, 2008).

¹² . A possible response here would be to insist that action is morally impermissible if it fails to express respect for the dignity of persons. This condition can be understood as imposing a formal or substantive condition. Understood formally, the proposal is unobjectionable, but does not impose substantive constraint. One expresses respect for the dignity of persons by treating them in whatever way moral theory says one ought to act. (If the correct moral theory is utilitarianism, then by treating people as utilitarianism dictates, one treats people as they ought to be treated, and expresses respect for the dignity of persons.) Understood as a substantive condition, the requirement of expressing respect for the dignity of persons is both unclear and controversial. The point is simply that we need to keep in mind the distinction between the two construals of the idea and not mix them together, appealing to the formal thought to show it is uncontroversially acceptable and then appealing to the substantive construal to show it has real content—against utilitarianism or some other form of consequentialism, say.

¹³ . A worry about claims that what makes a type of action morally wrong is the attitude it expresses is that the moral status of the act might seem to be part of what fixes the

message engaging in a type of action conveys. Theft, being morally wrong and a violation of the rights of the victim, might be thought to convey an attitude of disrespect towards its victim. Here at least the interpretation of what is expressed depends on a prior determination of wrongfulness. Also, if the message conveyed is what renders an act wrong, it seems one could always block the wrong by accompanying a doing of the act with an explicit sincere statement that one does not intend to convey or express the message that is standardly associated with this act.

¹⁴ . See Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge, MA: Harvard University Press, 1993).

¹⁵ . Kasper Lippert-Rasmussen, "Private Discrimination: A Prioritarian, Desert-Accommodating Account," *San Diego Law Review* 43 (2006), 817-856.

¹⁶ . Peter Railton, "Alienation, Consequentialism, and the Demands of Morality," *Philosophy and Public Affairs* (1982).

¹⁷ . Henry Sidgwick, *The Methods of Ethics* (New York: Dover Publications, 1966). This edition reproduces the seventh edition, published in 1907. First published 1884.

¹⁸ . See Derek Parfit, "What Makes Someone's Life Go Best," Appendix I in his *Reasons and Persons* (Oxford: Oxford University Press, 1984), 493-502.

¹⁹ . Care is needed in describing the implications of affirming welfarist act consequentialism that eschews interpersonal comparisons. These commitments only get you to the view described in the text if one holds nothing else matters. One might affirm further fairness norms; for the possibilities for the theory of justice that unfold on this terrain, see Marc Fleurbaey, *Fairness, Responsibility, and Welfare* (Oxford: Oxford University Press, 2008).

²⁰ . Richard H. McAdams, “Cooperation and Conflict: The Economics of Group Status Production and Race Discrimination,” *Harvard Law Review* 108, No. 5 (1995), 1003-1084.

²¹ . Elizabeth Anderson affirms something close to this democratic equality norm in her “What Is the Point of Equality?,” *Ethics* 109 (1999), 287-337. On the democratic equality ideal, see also Samuel Scheffler, “What Is Egalitarianism?” and “Choice Circumstance, and the Value of Equality,” reprinted in his *Equality and Tradition: Questions of Value in Moral and Political Theory* (Oxford: Oxford University Press, 2010). The third component of the ideal as stated in the text is of course a defining element of act consequentialism. Democratic equality incorporates the equality norm inherent in consequentialism but adds others.

²² . Samuel Scheffler, “*Choice, Circumstance, and the Value of Equality*,” section 7.

²³ . John Rawls affirms fair equality of opportunity in his *A Theory of Justice*, second edition (Cambridge, MA: Harvard University Press, 1999), ch. 2. Larry Temkin clearly affirms luck egalitarianism in his *Inequality* (New York and Oxford: Oxford University Press, 1993), at 13, to be read along with footnote 21 on the same page. There Temkin affirms that “it is bad—unjust and unfair—for some to be worse off than others through no fault [or choice] of their own. “

²⁴ . Robert Post describes the initially proposed Santa Cruz ordinance (which differs from the one the city eventually enacted) and ponders its implications for how we should conceive of our nondiscrimination practices in his “Prejudicial Appearances: The Logic of American Anti-Discrimination Law,” *California Law Review* 88 (2000), 1-40.

²⁵ . Sophia Moreau, “What Is Discrimination?”, *Philosophy and Public Affairs* 38 (2010), 143-179.

²⁶ . In the same vein, I do not see a principled basis for deciding what should count as the essence of a job for the purpose of deciding whether particular aspects of it provide a legitimate basis for a BFOQ exception to a disparate treatment claim, except by looking to the consequences of drawing these lines in one way rather than another. Is it acceptable to decline to hire men as nurses to care for elderly women who prefer to be tended by female nurses? Is it acceptable to decline to hire otherwise competent elderly persons or others who do not score high on a sexual attractiveness scale for jobs as sales staff in retail outlets that aim to sell cool, hip clothes to young people? On such questions, see Kimberly A. Yuracko, “Private Nurses and Playboy Bunnies: Explaining Permissible Sex Discrimination,” *California Law Review* 92 (2004); also Yuracko, “Trait Discrimination as Sex Discrimination: An Argument against Neutrality,” *Texas Law Review* 83 (2004).

²⁷ . “Because of” introduces some issues. One might be motivated not to hire Smith because he is black and one harbors animosity to blacks, when one also is motivated not to hire Smith because one correctly notices he is not qualified for the job. Must the former motivation be a “but for” cause of declining to hire Smith, in order that one should be found guilty of disparate treatment? Another issue is statistical discrimination. Being a member of a protected group may correlate with possession of traits relevant to being qualified for a job, such that a rational strategy for a firm considering applications is simply to eliminate from further consideration all members of the group in question.

Disparate treatment law forbids this sort of discrimination and requires a more individualized determination of the merits of applications.

²⁸ . See Joel Wm. Friedman, *Employment Discrimination: Examples and Explanations* (Austin and Boston: Wolters Kluwer, 2010).

²⁹ . The canonical presentation of Lockean libertarianism is in Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974). Lockean libertarianism falls outside what I have called the “society-wide consensus” on norms of nondiscrimination and the centrality of nondiscrimination norms to social justice.

³⁰ . Richard A. Epstein emphasizes this concern in his *Forbidden Grounds: The Case Against Employment Discrimination Laws* (Cambridge, MA: Harvard University Press, 1992). See also Epstein, “Should Antidiscrimination Laws Limit Freedom of Association? The Dangerous Allure of Human Rights Legislation,” *Social Philosophy and Policy* (2008), 123-156.

³¹ . I thank Thomas Hurka for very helpful comments on a draft of this essay.