

Chapter 4: Two Cheers for Affirmative Action

why supporters of racial preferences are wrong, and why their opponents are, too

A white candidate and a black candidate submit an application for a position. Both want to be admitted to the same university, hired at the same automobile plant, or promoted from their current jobs within the same police department. Both candidates are suitably qualified for the position: if either of them were the only one of the two to submit an application, the organization would offer the position to that candidate without hesitation. On paper, the white applicant is at least a little bit better qualified than the black applicant. His grades and test scores are a little bit higher, he has slightly stronger letters of recommendation, or he has a little bit more experience on the job. If the organization didn't know the race of either of the applicants, it would select the white applicant over the black applicant. But the organization does, in fact, know the race of both applicants. And it has a policy that leads it to select the black applicant over the white one, despite the fact that the white applicant is at least a little bit better qualified, precisely because the white applicant is white and the black applicant is black.

Supporters of policies that give at least some advantage to black applicants over white ones in this sort of way tend to refer to such policies as "affirmative action". It's hard to oppose taking action that is "affirmative," after all, without sounding, well, negative. Opponents of such policies, on the other hand, tend to refer to them as "racial preferences". If the advantage that's given to black candidates over white candidates is simply a matter of "preference" after all, rather than of reasoned judgment, then such policies seem to be simply a matter of imposing one's whim on others, like insisting that since you prefer chocolate ice cream to vanilla, everyone else should, too. I'm not particularly fond of either label. Something like "race-based prioritization" strikes me as a more accurate and less question-begging term to use. I'll mostly stick with the term "affirmative action" in this chapter, though, since that's the closest thing we have to an official title.

Affirmative action, in turn, is sometimes usefully divided into two categories: "weak" and "strong." Weak affirmative action involves taking positive steps to ensure that minority candidates are made aware of positions that are open and are encouraged to apply for them, but it then treats the applications that are received in an entirely race-neutral manner. In addition to doing whatever it would normally do to publicize a new opening, for example, a company that has a weak affirmative action policy might take out ads in magazines with a primarily black readership or on radio or television stations with a primarily black audience. It might send representatives to recruit at minority job fairs, make funding available to help black applicants fly in for interviews, and so on. But after doing all of this, a company that practices weak affirmative action would do nothing to give preference to black applicants over white applicants once all of the applications had been received. Potential black applicants would receive extra attention in the initial stage of the process, that is, when the company was trying to round up applications, but no extra attention in the final stage, when the applications it received were being evaluated. The people making the final decisions, for example, might not be given information about the race of the different candidates.

Strong affirmative action, on the other hand, goes beyond weak affirmative action by giving at least some sort of preference to minority candidates when the applications are actually being evaluated. In addition to assigning points to each applicant based on his grade

point average and SAT score, for example, a college that has a strong affirmative action policy might add some points to the total admission score of each black applicant. A black applicant might then end up with a higher total evaluation score than an otherwise comparable white applicant who had at least somewhat higher grades or test scores. Or the college might set aside a certain number of places in its freshman class for minority students and take its target numbers into account when examining its applicant pool. Either way, its policy of strong affirmative action would make it possible for a qualified black student to be admitted ahead of an otherwise better qualified white student precisely because the black student is black and the white student is white.

Weak affirmative action is generally uncontroversial. It simply involves making an effort to ensure that minority candidates are considered. Few people, if any, are opposed to it. But strong affirmative action is extremely controversial. My focus in this chapter will therefore be on it. Unless noted otherwise, when I use the expression “affirmative action” in what follows, I will really mean “strong affirmative action.” Strong affirmative action is controversial in a number of respects: people argue about its practicality, its morality, its legality and its Constitutionality. My focus here will be on the moral question posed by affirmative action policies, though I will at times make use of other sorts of arguments insofar as they help to clarify or advance some of the moral arguments. The central question that this chapter will address, then, is this: what is the moral status of strong affirmative action?

I have to admit that I’ve never found this to be a particularly interesting question. For as long as I can remember being aware of the existence of such affirmative action policies, I’ve always thought that the question had a pretty simple answer. And for as long as I can remember thinking this, I always thought that the answer was a pretty mundane, common sense, middle of the road one. I don’t remember ever reading anything about affirmative action or talking about it with anyone, but I guess I always assumed that most people thought that the correct answer was the answer that I always thought was the correct one. The answer itself didn’t strike me as particularly interesting or controversial, so I never really thought much about it. I certainly never contemplated writing about it.

When I first decided that I wanted to be able to teach a course on ethics and race, though, I began to look into the literature on the subject in some detail: books and articles by academics and media pundits, presentations by protestors and so on. And I discovered something surprising: it seemed that virtually no one was defending the view that I always thought was the sort of obvious view to hold. What I want to do in this chapter, then, is explain what my answer to the question is and how it differs from what seem to be the two much more common answers to the question, give a simple argument to establish a presumption in favor of my simple answer, and then respond to a number of objections that might be raised by defenders of either of the two more common positions. When I looked into the literature on slave reparations, I came to see that my initial view of that issue had been wrong. But the more I’ve looked into the literature on affirmative action, the more it’s come to seem to me that everyone else is wrong.

Let’s start by getting clear about what most people seem to think about affirmative action. Virtually everyone who writes or protests about it seems to fall into one or the other of two camps, what I’ll call the “pro” camp and the “anti” camp. Most people in the pro camp, those who support affirmative action, don’t just think that it’s nice to do black people a favor by adopting affirmative action policies. They think that having an affirmative action policy is a matter of justice, not a matter of charity. When Proposition 209 abolished affirmative action in California’s state university system, for example, these critics didn’t

simply complain that the state was eliminating something generous. They said that getting rid of affirmative action was positively immoral. Affirmative action, on this view, isn't simply something that's nice to do, like the campus sponsoring a free concert for the local community. Rather, affirmative action is something that it's positively wrong not to do.

Similarly, most people in the anti camp, those who oppose affirmative action, don't just think that affirmative action isn't such a great idea. They claim that it's positively unjust. If a state's universities practice affirmative action, for example, these critics don't simply say that the policy is foolish. They say that it's positively immoral. Affirmative action, on this view, isn't simply something that it might be unwise to do, the way that it might be unwise for the university's endowment to be invested in a portfolio that was too conservative or too risky. Rather, affirmative action is something that it's positively immoral to do.

Most critics of affirmative action think it's downright wrong to practice affirmative action, then, and most defenders of affirmative action think it's downright wrong not to. As a result, people in both groups are often easily agitated and quick to accuse people in the other group of racism (or "reverse racism"). Those who oppose affirmative action policies are often angry when such policies are applied to them. Those who support affirmative action policies are often angry at those who are angry at affirmative action. Both groups are quick to quote Martin Luther King, Jr. and to claim that if he were alive today, he would agree with them.¹ And both groups are prone to write or speak as if the debate over affirmative action can only have these two sides. One popular anthology on the subject, for example, is called *Affirmative Action: Social Justice or Reverse Discrimination?*² This title reflects the widespread assumption that affirmative action must either be demanded by justice or be prohibited by it. The assumption that there's only room for these two sides in the affirmative action debate, moreover, is also frequently made explicit in the literature on the subject. The philosopher John Kekes, for example, introduces an article of his on the issue by saying that it sets out "to advance the claim that one side has it right and the other is mistaken."³ The philosopher Alan Goldman writes that a middle ground on the subject is so hard to imagine that "Interested persons coming to this literature for the first time may have no choice but to review its now extensive history and make an informed choice to come down on one side or the other."⁴ And in the opening paragraph of his recent book *Understanding Affirmative Action*, which is designed to provide a balanced and up-to-date introduction to the subject, Public Administration and Policy Professor J. Edward Kellough writes that "People disagree on whether affirmative action should be permitted or, if it is judged to be necessary, on the specific types of efforts that should be included," as if the only options are that affirmative action is impermissible or necessary.⁵

¹ Critics of affirmative action almost universally appeal to the line in King's famous "I Have a Dream" speech in which he dreams that his children "will one day live in a nation where they will not be judged by the color of their skin but by the content of their character." On the other hand, King also wrote the following: "Whenever this issue of compensatory or preferential treatment for the Negro is raised some of our friends recoil in horror. The Negro should be granted equality, they agree; but he should ask for nothing more. On the surface, this appears reasonable, but it is not realistic. For it is obvious that if a man is entering the starting line in a race three hundred years after another man, the first would have to perform some impossible feat in order to catch up with his fellow runner" (quoted in Moore 2005: 88).

² Beckwith and Jones (1997).

³ Kekes (1993: 144).

⁴ Goldman (1993: 299).

⁵ Kellough (2006: 3).

But this common assumption that affirmative action must be either morally required or morally prohibited is a mistake. Some things are forbidden by justice and some things are required by justice, but plenty of other things fall in between. Justice doesn't require you to put your left shoe on before your right, for example, but it doesn't forbid you from doing so either. Unlike most people who write about affirmative action these days, and unlike virtually everyone who protests for or against it, I think that affirmative action falls into this in between area.⁶ So far as morality and justice are concerned, that is, schools and businesses and governments are free to practice affirmative action if they wish and free to decline to practice it if they wish. If a school, or business, or government practices affirmative action, that's not morally wrong. If a school, or business, or government doesn't practice affirmative action, that's not morally wrong either. I'll try to explain why I think that's true in the section that follows, and I'll try to respond to a number of objections to my view in the sections that follow that.

It's possible, of course, that in trying to stake out a middle ground between the view that practicing affirmative action is immoral and the view that not practicing it is immoral, I'm misunderstanding the extent to which my own view is out of step with that of most people who write and talk about this issue. Perhaps most people who oppose affirmative action don't really mean it when they use words like "unjust" or "immoral" to describe the practice or when they complain that affirmative action violates the "rights" of white applicants. Perhaps these are just rhetorical flourishes that disguise the more mundane view that affirmative action, although not morally wrong, isn't a good idea. And perhaps most people who support affirmative action don't really mean it when they use words like "unjust" or "racist" to refer to initiatives like Proposition 209 that prohibit affirmative action or when they call affirmative action itself "morally mandatory."⁷ Perhaps these, too, are just rhetorical flourishes, obscuring the more mundane view that even though it's not immoral to abolish affirmative action, it would be better not to abolish it. If most people already believe that it's not wrong to practice affirmative action and that it's not wrong not to practice it, then it will turn out that the goal of this chapter is to provide a justification for what most people already believe, rather than to provide a moderate alternative to what most people believe. And since the goal of this book is to try to answer the question of whether race should matter in a variety of contexts by developing positions that can be justified on grounds that most people accept, rather than to engage in contrarian argumentation for its own sake, I'll be more than happy if it turns out that I'm simply making explicit what most people are already tacitly thinking and then providing a careful and detailed defense of it.

I'm generally inclined to take people at their word, though, and to take the arguments they present at face value. And on the face of it, at least, most critics really do seem to be saying and setting out to show that affirmative action is positively immoral, and most

⁶ For examples of supporters of affirmative action who believe that affirmative action is required by morality, see Armour (1997), Moore (2005), Thomas (1993), Francis (1993), Turnbull (1993), Arneson (who calls it "morally mandatory") (1993: 164), and Sterba (1994). For examples of opponents of affirmative action who believe that affirmative action is prohibited by morality (typically calling it "unjust" or "immoral"), see Kekes (1993), Simon (1993), Hartle (1993), Quinn (1993), Olafson (1993: 211) and Sommers (1993). It's hard to find anyone who has explicitly tried to stake out the kind of middle ground I seek to occupy in this chapter, but one exception worth reading is Thomas Nagel's (1973) article. In this paper, Nagel argues that affirmative action is not required by justice but also argues that it is not "seriously unjust" either. But while the position he is led to is thus one of the few to come close to the one I try to defend here (I say close because Nagel leaves it open that affirmative action might still be moderately unjust, and I don't think that's true, either), the argument he presents in its defense is quite different from the sort of arguments I develop in this chapter.

⁷ Arneson calls it "morally mandatory") (1993: 164),

supporters really do seem to be claiming and trying to show that abolishing affirmative action is positively immoral. If most of these people don't really mean what they seem to be saying, it's somewhat difficult to understand what they're getting so worked up about. If you concede that affirmative action isn't really immoral and doesn't really violate anyone's rights, after all, you might still think it's a bad idea but why would you try to ban it? If you admit that no one has a right to the benefits of affirmative action and that abolishing affirmative action isn't immoral, you might still think that getting rid of it's a mistake, but why would you accuse those who seek to do so of being immoral or racist? In any event, regardless of whether it's common to believe that neither practicing nor abolishing affirmative action is immoral, that is the view that I will attempt to defend in this chapter.

confessions of a white affirmative action baby

One of the most widely quoted works on the subject of affirmative action is Stephen Carter's thoughtful book, *Confessions of An Affirmative Action Baby*. In it, Carter discusses the fact that he was admitted to the Yale Law School under an affirmative action program that gave priority to qualified black applicants like himself over white candidates who had even higher grades and test scores. Describing his admission under such circumstances as a "mixed blessing," Carter poignantly depicts the various psychological costs that can burden those who are meant to benefit from affirmative action policies.⁸ Anyone who is seriously concerned with the issue of affirmative action in America should certainly read and think about Carter's book.

I have something in common with Stephen Carter. I got into Yale under an affirmative action policy, too. The policy that I benefited from helped me get into Yale College rather than Yale Law School, and it benefited me because of my geographical identity rather than because of my racial identity. But I'm inclined to think that the similarities between the two cases are far more important than the differences. Indeed, the more deeply I've looked into the literature on race-based affirmative action, the more firmly I've come to believe that the similarities between geography-based and race-based affirmative action can help point the way to a resolution of the debate over race-based affirmative action, a kind of middle ground between the pro and the anti camps that appeals to considerations that people in both groups already accept.

Let me start by explaining what I mean by geography-based affirmative action. When a student applies to a school like Yale, the admissions office puts a good deal of weight on things like the student's grades, test scores, extra-curricular activities, teacher recommendations, and so forth. But schools like Yale typically care about more than this. They want, among other things, to ensure that their students are exposed to people who come from backgrounds that are different from their own. One way to do this is to try to bring together students from as many parts of the country as possible. And one way to make this happen is to give some preference to applicants from parts of the country that are underrepresented on campus. If there are hardly any students from North Dakota in the current student body, for example, and if there are already a very large number of students from New York, then an applicant with very good grades and test scores from North Dakota might well be admitted ahead of an applicant with somewhat better grades and test scores from New York. Like Stephen Carter, I can't be sure that affirmative action helped make the

⁸Carter (1991: 71).

difference in my admission to Yale. But I do know this: I was admitted to every college I applied to on the East Coast and was not admitted to the one school I applied to on the West Coast. Having grown up in a Western state, geography-based affirmative action counted in my favor when I applied to schools on the East Coast and counted against me when I applied to schools on the West Coast. This is the reality of geography-based affirmative action.

Setting aside the issue of race-based affirmative action for just a moment, let's consider the question of what we should say about the moral status of geography-based affirmative action. I think that most people, regardless of what they would say about race-based affirmative action, would agree about two things. First, morality allows organizations to practice geography-based affirmative action. To put things in slightly more philosophical terms, geography-based affirmative action is morally permissible. If Yale admits a slightly less qualified student from North Dakota over a slightly better qualified student from New York, that is, the student from New York can't really complain that his moral rights have been violated. So far as I can remember, I've never heard anyone claim that a policy such as Yale's is positively unjust.⁹ Nor, for that matter, have I heard anyone complain about the geography-based affirmative action policies employed by the United States Military Academies. In order to be admitted to West Point, for example, students must first secure a nomination to be considered, and most of the nominations that are available are distributed by congressional district. This makes it considerably easier to get into West Point if you grow up in some parts of the country than if you grow up in others. But, again, everyone seems to agree that, morally speaking, West Point has the right to make its admissions decisions in this way if it thinks it important that future military leaders be drawn from, or learn to work alongside people who are drawn from, all parts of the country. So geography-based affirmative action, whether practiced by a private organization or by an agency of the government, strikes nearly everyone as morally permissible. It's not forbidden by considerations of morality or justice.

The second thing I think most people would agree about with respect to geography-based affirmative action is that it isn't required by considerations of morality or justice, either. Organizations have no moral duty to engage in it or, again to put things a bit more philosophically, geography-based affirmative action is non-obligatory. If Yale or West Point decided to eliminate its geography-based affirmative action policy, that is, and if it therefore decided to admit a slightly better qualified student from New York over a slightly less qualified student from North Dakota, the student from North Dakota couldn't really complain that his moral rights had been violated, either. Neither the slightly better qualified student from New York nor the slightly less qualified student from North Dakota has a moral right to be admitted to West Point or Yale ahead of the other. Morally speaking, the schools are free

⁹ One possible exception that is worth noting could arise in cases where a school adopts a geography-based affirmative action program in order to discriminate against a disliked racial or ethnic group. And there is, in fact, some evidence that elite universities like Harvard, Yale and Princeton originally adopted such policies for precisely this reason. The anti-Semitism then prevalent at such institutions led them to want to limit the number of Jewish students they admitted, and one way to do this was to limit the number of applicants accepted from cities that had large Jewish populations (see, e.g., Arthur (2007: 239) and Karabel (2005)). It seems plausible to me to think that if an otherwise morally acceptable policy is adopted in order to achieve a morally objectionable goal, then the act of adopting the policy for that reason is morally objectionable even though the policy in and of itself is not. Even if it can be wrong to adopt a geography-based affirmative action program in this way, however, this will have no bearing on the discussion of race-based affirmative action in this chapter. No one really thinks that affirmative action policies that favor black applicants over white applicants are adopted out of contempt for white people, and so while there might be a reason to think that some cases of geography-based affirmative are morally objectionable, there will be no comparable reason to think that any cases of race-based affirmative action are.

to admit whichever of the two qualified candidates they prefer. Geography-based affirmative action, in short, is permissible but non-obligatory. It's morally optional.

I think that virtually everyone, regardless of their view of race-based affirmative action, recognizes that geography-based affirmative action is permissible but non-obligatory. This is why it's been possible for geography-based affirmative action to be employed by many schools for as long as there have been formal admissions policies without generating any real controversy and why it's also been possible for many other schools not to practice geography-based affirmative action without generating any real controversy either.¹⁰ If an organization practices it, that's fine. If an organization doesn't practice it, that's fine, too. Either way, virtually no one complains about it or files lawsuits about it.

But geography-based affirmative action seems to have the morally significant features that people on both sides of the debate over race-based affirmative action typically point to when they argue for or against it. Defenders of race-based affirmative action, for example, typically appeal to the fact that some races have been historically disadvantaged relative to others, that some races are currently disadvantaged relative to others, or that adding members of underrepresented races to an organization can benefit it by increasing the diversity of perspectives it can make use of. All of these things are also true of geography-based affirmative action: some parts of the country have long been disadvantaged relative to others, some parts of the country are currently disadvantaged relative to others, and schools where all the students are from the same area provide a less stimulating environment than schools whose students come from all across the country. Opponents of race-based affirmative action typically point out that you can't control what race you're born into, that generalizations about racial groups aren't true of each individual member of those groups, and that race-based affirmative action takes group membership into account rather than judging people solely on their individual merits. All of these things are true of geography-based affirmative action, too: you can't control what part of the country you're born and raised in, generalizations about people from different parts of the country aren't true of each individual in those parts, and geography-based affirmative action takes group membership into account rather than judging people solely on their individual merits. Race-based affirmative action, in short, seems to be morally on a par with geography-based affirmative action. What morality says about one it seems it should also say about the other.¹¹

¹⁰ Dartmouth became the first college to adopt a formal admissions policy in 1922, and one of its nine criteria for admissions was "the principle of geographical distribution". And Harvard can trace its geography-based affirmative action practices back even further, at least as far back as the tenure of President Charles W. Eliot, who presided over the University from 1896 to 1909. Eliot's approach to student admissions put weight on "the wholesome influence that comes from observation of and contact with" people from differing backgrounds and thus sought to bring to Harvard students "from North and South, from East and West" (quotes taken from Fullinwider and Lichtenberg (2004: 83-4, 165, 166)).

¹¹ It might at first seem that there's an important difference between the two cases insofar as you can change your geographical location but not your race, but this appearance is misleading: a student whose parents choose to raise him in New York rather than in North Dakota can no more make it the case that he is raised in North Dakota than a student whose parents are white can make it the case that he is black. It's true, of course, that some students may credibly be able to present themselves as residing in more than one state. If a student's parents are divorced and have joint legal custody of him and one of them lives in New York and the other lives in North Dakota, for example, then the student may be able to choose which state to identify himself with when he applies to college. But, then again, if a student's parents are of sufficiently mixed ancestry, he, too, may find that he can try to "pass" as either black or white depending on the circumstances; there is no real difference between the cases here.

These considerations suggest a simple two-step argument in support of the view that race-based affirmative action is permissible but non-obligatory. In fact, they suggest two such arguments. One argument takes as its first step the claim that geography-based affirmative action is permissible but non-obligatory and adds as its second step the claim that race-based affirmative action really is morally on a par with geography-based affirmative action. If both steps of this argument are accepted, it follows that race-based affirmative action is permissible and non-obligatory, too. The second argument also takes as its first step the claim that geography-based affirmative action is permissible but non-obligatory but it adds as its second step the more modest claim that we have *some reason* to think that race-based affirmative action is morally on a par with geography-based affirmative action. If both steps of this second argument are accepted, it doesn't necessarily follow that race-based affirmative action really is permissible and non-obligatory. It simply follows that we have some reason to think that it is.

After having spent a good deal of time trying to work my way through the literature on affirmative action, I find myself inclined to accept the stronger of these two arguments. Although there are clearly differences between race-based and geography-based affirmative action, and although in the case of at least some of the differences I can see the initial force behind the thought that the difference in question is a morally significant one, in the end I've come to think that none of the things that distinguish race-based from geography-based affirmative action are sufficient to render its race-based form impermissible if its geography-based form is permissible or to render its race-based form obligatory if its geography-based form is non-obligatory.

But while the stronger version of the argument seems right to me, I don't expect it to seem right to most readers, at least not right away, because I don't expect the stronger version of the second step to seem right without first considering a number of objections that might be made against it. For the purposes of this chapter, then, I simply want to appeal to the weaker version of the argument. The second step of the weaker version merely maintains that we have *some reason* to think that the two forms of affirmative action are morally on a par, and I think I've already shown that we have some reason to think this by showing that the morally relevant features that appear most conspicuously in arguments for and against race-based affirmative action are also features of geography-based affirmative action. The conclusion of this more modest argument isn't that race-based affirmative action is permissible and non-obligatory, but simply that it's reasonable to start with an initial presumption in favor of the view that race-based affirmative action is permissible and non-obligatory. With that presumption established, then, the question is whether there are any arguments on either side of the issue that are strong enough to overcome it. I think that there aren't and that's what I'll try to show over the course of this chapter.

Before turning to the details of the various arguments that can be given for and against affirmative action, though, it's worth noting one possible source of confusion that may arise at this point. For it may seem that I've already insisted that there really is an argument strong enough to overcome this presumption and that, in particular, I've already committed myself to the view that affirmative action, far from being morally optional, is really morally obligatory. In chapters 2 and 3, after all, I presented and defended an argument in favor of slave reparations. The reparations position maintains that African Americans are entitled to be benefited today because of the wrongful harms that were inflicted on Africans and African Americans in the past. Race-based affirmative action is designed to benefit black Americans. Since race-based affirmative action benefits black

Americans and since the reparations position that I've already defended maintains that black Americans are entitled to be benefited, it may seem that I've already argued in defense of the claim that black Americans are entitled to race-based affirmative action. If black Americans are entitled to race-based affirmative action, of course, then race-based affirmative action is morally obligatory, not morally optional. And so for me to argue in this chapter that race-based affirmative action is not, in fact, morally obligatory (or morally impermissible) might seem inconsistent with what I myself have maintained in the previous two chapters.

While this diagnosis of my position is understandable, and while most people who support slave reparations probably do also support affirmative action, there is in fact nothing inconsistent about agreeing that reparations are obligatory while denying that affirmative action is obligatory. There's nothing inconsistent, for that matter, about agreeing that reparations are obligatory while denying that affirmative action is even morally permissible. Saying that the government has an obligation to provide some sort of benefit to a group of people, after all, isn't the same as saying that the government has an obligation to provide every kind of benefit that it could provide to that group, nor is it the same as saying that it has an obligation to provide any one particular kind of benefit to it. Providing free toothpaste to every black American would clearly benefit them, for example, but believing in the reparations position doesn't entail believing in an obligation to provide free toothpaste to all black Americans. Stealing money from the Queen of England and distributing it to black Americans would clearly benefit them, too, but believing in the reparations position doesn't entail that doing so would be morally permissible, let alone morally obligatory. The fact that affirmative action is one kind of benefit that black Americans could be offered, then, isn't enough to show that an obligation to provide reparations for slavery and its aftermath entails an obligation to engage in race-based affirmative action. There's nothing inconsistent about believing that there's an obligation to provide some kind of benefit without believing that there's an obligation (or even a right) to provide this particular kind of benefit.

In other sorts of cases in which most people agree that a government owes compensation to a group of people, moreover, virtually no one thinks this means that it owes the benefit of affirmative action to that group in particular. Most people who agree that the United States owed reparations to Japanese-Americans who were wrongfully interned during the second world war, for example, don't think that this means that they were owed affirmative action in particular, nor do those who agree that the German government owed reparations to victims and families of victims of the Holocaust think this means that the German government was obligated to provide them with affirmative action benefits, say by giving Holocaust survivors preferential treatment in applying to work or study in German universities. And the same is true in other sorts of cases in which most people agree that government compensation is owed. If the government wrongfully buries toxic waste in certain neighborhoods, or wrongfully frames and imprisons people who are guilty of no crime, most people will agree that the government would owe them some sort of reparations as a result but, again, virtually no one would insist that it owes them the benefits of affirmative action in particular, say in being considered for admission to West Point or for a job with the postal service. From the fact that a government owes a group of people reparations for harm that it wrongfully caused them, it again doesn't follow that it owes them the benefits of affirmative action in particular.

None of this, of course, is enough to show that affirmative action isn't morally obligatory. But it should be enough to show, at the very least, that a commitment to the reparations position isn't enough by itself to justify a commitment to the claim that

affirmative action is morally obligatory. There's nothing inconsistent about maintaining that the former is obligatory while maintaining that the latter isn't. And so whether one accepts or rejects the arguments that I offered in chapters 2 and 3, we are left in the same position with respect to the question of the moral status of affirmative action that we started with: the case of geography-based affirmative action gives us at least some reason to start with a presumption in favor of the view that race-based affirmative action is permissible but non-obligatory, and we must now consider the question of whether there are any arguments strong enough to overcome this initial presumption. I haven't yet tried to show that there aren't any. But I've at least tried to show that accepting the slave reparations position doesn't mean that there are.

objections from the left

My claim about race-based affirmative action is a moderate one. It says that morally speaking, organizations may practice affirmative action if they wish to, but that morally speaking, they aren't required to do so if they don't. Because my position attempts to occupy a kind of middle ground in the debate over affirmative action, it can be attacked from two very different sides. Those who support affirmative action will welcome my claim that affirmative action is permissible, but most of them will reject my claim that it isn't obligatory. Those who oppose affirmative action will welcome my claim that it isn't obligatory, but most will reject my claim that it's permissible. I'll try to address these two sets of complaints in turn, beginning with the arguments most commonly presented by defenders of affirmative action. Some of the arguments that have been offered in defense of affirmative action explicitly aim to show that affirmative action is morally obligatory. Others are somewhat ambiguous between claiming to show that affirmative action is morally obligatory and merely claiming to show that it's a good (but morally optional) thing to do. And still others, perhaps, are offered as an attempt to identify something good about affirmative action without really trying to show that it's wrong not to practice it at all. Whatever the intentions behind the various arguments that have been made in defense of affirmative action have been, though, they pose a challenge to my claim that affirmative action is permissible and non-obligatory only if they can do something to show that it's positively immoral not to practice affirmative action. For the purposes of the discussion that follows, then, I'll focus on the question of whether any of them can.

objection one: the unfair disadvantage argument

One kind of argument in favor of affirmative action is essentially backwards looking. This kind of argument points to the many harmful injustices that were committed against black Americans in the past and claims that the lingering effects of these historical wrongs can be used to ground an obligation to use affirmative action as a means of benefiting black Americans in the present. This is the kind of justification that defenders of affirmative action originally provided for such policies in the 1960s.

The most familiar version of this backwards-looking kind of argument appeals to the idea of fairness. The argument is often motivated by appealing to the example of runners competing in a race, and it goes back at least as far as an early speech on the subject by President Lyndon Johnson, who famously declared in 1965 that "You do not take a person who, for years, has been hobbled by chains and liberate him, bring him up to the starting line of a race and then say, 'You are free to compete with all others,' and still just believe that you

have been completely fair.”¹² This appeal to fairness as a justification for affirmative action has a good deal of intuitive force behind it, and the argument based on it has been picked up on more recently by a number of people. The philosopher James Sterba, for example, defends affirmative action for those candidates who “have the potential to be as qualified or more qualified than their peers, but [whose] potential has not yet been actualized because of past discrimination and prejudice.” His claim is that it isn’t fair to reject a candidate who would have won a position had he not been hampered by the effects of previous injustices. And, echoing Johnson’s famous words, he attempts to support his position by appealing to the case of a runner who is unfairly hampered in a race: “persons who receive preferential treatment are like runners in a race who for a time are forced to compete at a disadvantage with other runners, e.g., by having weights tied to their legs, but then after are allowed to transfer those weights to the runners in the race who had previously benefited from the unfair competitive advantage so that the results of the race will now be fair.”¹³

This unfair disadvantage argument for affirmative action depends on two claims. The first is the claim that black Americans who apply for jobs and admission to schools in the present are unfairly disadvantaged by the effects of unjust discrimination that took place in the past. White Americans today enjoy an “unfair competitive advantage,” in other words, because of the misdeeds of the white Americans of yesterday. The second is the claim that if this first claim is true, then it would be morally wrong for those who evaluate such applications to fail to adjust for this fact in making their final decisions about which applications to accept. Knowing that some entrants in a competition have been unfairly disadvantaged, in other words, those who sponsor the competition are morally obligated to take this fact into account before deciding on the winners. Perhaps not surprisingly, I think that the first claim made by the unfair disadvantage argument should be accepted. The first claim, after all, is basically the same as the third step in the compensation argument for slave reparations that I endorsed in chapters 2 and 3, and I defended the claim in that context at some length. But I think that the second claim made by the unfair disadvantage argument should be rejected. And that’s why I think that the unfair disadvantage argument itself should be rejected.

To see what’s wrong with this part of the unfair disadvantage argument, it helps to begin by looking at a simpler case that focuses on just a single pair of applicants and that doesn’t involve race. So consider the story of Bill and Ted, two white students who go to the same high school. During their freshman and sophomore years, Bill got somewhat better grades than Ted and did somewhat better on practice SAT tests. There was thus every reason to believe that Bill was going to be somewhat more qualified and better prepared for college than Ted was going to be. Over the summer between their sophomore and junior year, though, Bill’s father was unjustly framed and convicted for a crime that he didn’t commit. As a result, Bill’s life was seriously disrupted. His father’s being sent to prison was emotionally painful for Bill. Bill was forced to get a job working after school. He had less time to spend on his homework and even when he worked on it he had more difficulty concentrating. His study habits declined somewhat because he no longer had a father at

¹²Quoted in Moore (2005: 80).

¹³Sterba (1993: 287). Sterba defends this approach with a caveat – only in cases where the winner can, in effect, be expected to catch up with some extra help (1993: 286-7). Sterba also defends this kind of argument: “the key to an adequate justification” of affirmative action, he says, is “to see that practice, not as the redressing of past privations, but rather as a way of neutralizing the present competitive disadvantage caused by those past privations and thus as a way of restoring equal access to those goods which society distributes competitively” (1975: 53).

home to keep after him, his grades went down a bit as a result, and he didn't end up doing as well on the SAT as his earlier practice scores would have predicted. By the time that college applications came around, in fact, Ted had a somewhat better overall academic record than Bill and was somewhat better prepared to succeed in college.

Now let's suppose that you're the admissions officer at a college that Bill and Ted have both applied to. Both Bill and Ted are suitably qualified for admissions to your school. Ted's academic preparation and record are somewhat better than Bill's. But you know from reading their letters of recommendation that Bill's would almost certainly have been somewhat better than Ted's had Bill's father not been unjustly convicted of a crime he didn't commit. You only have room to admit one more student. And now ask yourself this question: is it okay for you simply to admit Ted because he is the better qualified candidate or must you instead take into account and adjust for the fact that Bill has been unfairly disadvantaged by the wrong that was done to his father?

Notice that the question is not whether you *should* simply admit Ted. I think a reasonable case could be made for admitting either candidate. The question is whether it would be *okay* for you simply to admit the white student who in fact has the better academic preparation and record rather than give preferential treatment to the white student who almost certainly would have had the better academic preparation and record had the world been a more fair place. And here, at least, the answer seems clear: you have every right to admit Ted to your school over Bill. Although it would be considerate of you to take Bill's misfortune into account, you do nothing morally wrong if you decide simply to admit the white student who is, in fact, better qualified for admission over the white student who is, in fact, less qualified. I suspect that most people, regardless of their views about race-based affirmative action, would accept this judgment.

But if I'm right about this,¹⁴ an important conclusion follows, and this conclusion undermines the second claim made by the unfair disadvantage argument. For it's clear in the case of Bill and Ted that Ted enjoyed an "unfair competitive advantage" over Bill. Ted was able to spend more time on his homework than Bill because he didn't have to get a job after school, he was able to concentrate more when he did his homework than Bill because he had a more stable family life than Bill had, and he continued to reap the benefits of having a father in the house to help keep him in line, an asset of whose value Bill never fully appreciated until he lost it. All of this is true because Bill's father, but not Ted's, was unjustly convicted of a crime that he didn't commit. And this is clearly unfair. But while it's clear that Bill therefore suffers an unfair disadvantage relative to Ted, it's equally clear that morality permits a school to admit Ted over Bill without adjusting for this fact. And this can only mean one thing: the schools that Bill and Ted apply to are not morally obligated to

¹⁴ It may be worth noting that even if someone insists that it would be positively immoral for you to admit Ted rather than Bill, this response is unlikely to lend support to the claim that race-based affirmative action is morally obligatory. If you think that it's unfair for Bill to be penalized for the fact that his father went to jail, after all, then you'll probably think this regardless of whether or not Bill's father was guilty. Either way, after all, the fact that his father went to jail was not Bill's fault, and so, either way, you'll think that the university would be obligated to take this fact into account. You'll think that it's wrong to let Ted's financial advantage over Bill help Ted get admitted ahead of Bill not because Ted's financial advantage arose from some unjust action that took place in the past, that is, but simply because you'll think that it's wrong for any financial advantage to help one student get admitted ahead of the other. While this kind of reasoning might therefore help to defend some form of class-based affirmative action, then, it will not help to justify an obligation to engage in race-based affirmative action in particular, and my focus in this chapter is exclusively on race-based affirmative action.

adjust for the fact that Bill suffers an unfair disadvantage relative to Ted in making their final decision about which application to accept. Knowing that some entrants in the competition for admission to their school have been unfairly disadvantaged, in other words, they are not morally obligated to take this fact into account before deciding on the winners. The second claim needed to sustain the unfair disadvantage argument is therefore false, and shown to be false, on terms that virtually everyone on both sides of the debate over affirmative action will accept.

And, indeed, when the attempt to look at affirmative action through the lens of an unfair race is carried out with this point in mind, the same lesson about morality and fairness emerges from the very sort of story that defenders of affirmative action have so frequently appealed to. For suppose, instead, that Bill and Ted had each been preparing to try out for the last available slot on the track team. In previous years, they trained for the same amount of time each week and Bill had been a little bit faster and more consistent than Ted. After Bill's father was sent to prison and Bill had to get a job after school, though, Bill gained weight from overeating as a coping mechanism, had less time to practice running because of his job, and missed the firm but loving discipline his father had always provided. As a result of this, by the time the day of the tryouts arrived, Ted was a little bit faster and more consistent than Bill. Virtually everyone, regardless of their view of race-based affirmative action, will agree that the track coach would do nothing morally wrong in this case if he simply gave the slot on the team to the runner who is in fact faster and more consistent rather than adjusting for the fact that a different runner would instead have been a bit faster and more consistent had he not been unfairly disadvantaged by the unjust harm that was inflicted on his father. And so in this case, again, it should be clear that it's perfectly acceptable simply to select the person who is in fact better qualified over the person who would have been better qualified if the world has been a fairer place.

In cases where two white students compete for a position, then, it should be clear that the mere fact that one is better qualified than the other only because the other has been unfairly disadvantaged doesn't mean that it's wrong simply to accept the more qualified applicant without taking this unfortunate fact into account. But if this is true in cases where both applicants are white, then it must be equally true in cases where one is white and the other is black: either the existence of an unfair disadvantage generates a moral obligation to take the disadvantage into account or it doesn't. And since the existence of such a disadvantage isn't enough to generate such an obligation in the case where both candidates are white, the existence of such a disadvantage can't be enough to generate such an obligation in the case where one of the candidates is white and the other is black. The fact (if it is a fact) that a given black applicant would have been better qualified than a given white applicant had the black applicant not been unfairly disadvantaged by the legacy of slavery and its aftermath, then, can't, by itself, show that it would be wrong simply to admit the white applicant over the black applicant without taking these circumstances into account.

The fact that it's okay not to give Bill preferential consideration relative to Ted, of course, doesn't by itself mean that it's okay not to give black candidates preferential consideration relative to somewhat better qualified white candidates. There are a number of other features of race-based affirmative action that might be used to try to justify it, and many of these might not apply to the case of Bill and Ted. But the unfair disadvantage argument appeals to one feature of race-based affirmative action in particular: the existence of an unfair disadvantage between the black candidates and the white candidates. The case of Bill and Ted is enough to show that this particular feature isn't powerful enough to generate an

obligation to provide preferential consideration, and this, in turn, is enough to show that the unfair disadvantage argument isn't powerful enough to generate an obligation to engage in race-based affirmative action. It may be considerate or kind or even commendable or praiseworthy to take Bill's circumstances into account in the case of Bill and Ted, but it would not be positively immoral to decline to do so. And so while the unfair disadvantage argument may show that there's something good about having a race-based affirmative action policy, it can't show that there's anything positively immoral about not having one.

There is a second kind of backwards-looking argument that also merits consideration, but before I say something about it, I want to acknowledge some objections that might be raised at this point against my use of the story of Bill and Ted in arguing against the unfair disadvantage argument. For there are clearly some noticeable differences between the unfair disadvantage involved in the case of Bill and Ted, on the one hand, and unfair disadvantage involved in the case of race-based affirmative action, on the other. And it might well be thought that these differences undermine the objection to the unfair disadvantage argument that I've raised here. If the differences turn out to be morally relevant, after all, then they might explain how it could be morally acceptable to ignore the unfair disadvantage and select Ted over Bill but not morally acceptable to have a policy of ignoring the unfair disadvantage and admitting slightly better qualified white applicants over slightly less qualified black applicants. There are four differences, in particular, that a careful reader may well be wondering about at this point.

One difference is simply this: the case of Bill and Ted involves a single, isolated incident. But there are a great many cases in which a slightly less qualified black candidate has been unfairly disadvantaged relative to his slightly better qualified white rival. It might be thought that this difference makes a difference.

This concern strikes me as unwarranted for two reasons. First, it's hard to see why the numbers involved should be morally relevant. In trying to figure out whether it's okay to select one particular candidate over another candidate who's a bit less qualified as a result of some unfair disadvantage, why should it matter how many other people find themselves in a similar situation? Why would it be okay to simply select the better qualified candidate if the situation is rare but not okay to do so if the situation is common? Why should the merits of Bill's plea to have his family's misfortune taken into account be made to depend on whether or not other families have suffered a similar injustice? Second, this suspicion that the difference in numbers is irrelevant seems to be vindicated if we change the case of Bill and Ted to make it more like the case of race-based affirmative action in terms of the number of people involved. Suppose, for example, that I add that Bill isn't the only white person in the country to have been hindered by his father having been wrongfully convicted of a crime he did not commit. Suppose it turns out that there are a great many such white people, all of whom find themselves a bit less qualified for a position they apply for than some other white candidate who has not suffered a similar misfortune. Does the existence of this large number of other unfairly disadvantaged white candidates change your intuition about the case of Bill and Ted? I, at least, find that it doesn't. It seems just as clear to me in this version of the story as in my initial version that it's perfectly acceptable simply to select Ted over Bill even though Bill would have been better qualified than Ted had Bill's father not been unjustly imprisoned. It's hard to believe that many people on either side of the debate over race-based affirmative action would disagree. And so it's hard to see how this first objection to my argument against the unfair disadvantage argument could be accepted.

A second difference between the two cases might seem to pose a bigger problem for my argument. The point isn't just that there are a large number of unfairly disadvantaged black candidates while there's only one unfairly disadvantaged Bill, a critic of my position might say, but rather that the black candidates have been disadvantaged systematically, while Bill has been disadvantaged as the result of a more or less random act of wrongdoing. Indeed, my impression from talking to a number of people who support affirmative action is that this is likely to be one of the biggest complaints made about my argument to this point: that my appeal to the intuitive reaction to individual cases, even if we stipulate that there are a great many of them, obscures the systematic nature of racial disadvantage in this country.

While I realize that many people are inclined to put a great deal of weight on this point, however, I have to admit that I have difficulty seeing what reason could be given for doing so. Indeed, the same two considerations that lead me to reject the first objection to my argument lead me to reject this second objection as well. First, at a general level, it just isn't clear why the difference between being a victim of systematic harm and being a victim of non-systematic harm should be considered morally relevant. In trying to figure out whether it's okay to select one particular candidate over another candidate who is a bit less qualified as a result of some unfair disadvantage, why should it matter how the unfair disadvantage was caused? It's not hard to understand why it might matter whether the disadvantage was caused fairly or unfairly. If one candidate is disadvantaged as a result of some harm that was caused fairly, after all, then presumably no one will think that we're obligated to select him over another better qualified candidate. If Bill had been forced to work long hours after school instead of doing his homework or training for the track tryouts because of some event that was clearly his own fault, for example, then it's not hard to see how that would make things importantly different from the case of black candidates who are disadvantaged due to no fault of their own. But given that the disadvantage that Bill suffers in the case that I described is just as unfair to him as the disadvantage that unfairly disadvantaged black candidates suffer is unfair to them, why should it matter that the unfair disadvantage is caused systematically in the latter case and non-systematically in the former? Why would it be okay to select the better qualified candidate over the one who would have been better qualified had he not been unfairly disadvantaged by a non-systematically caused harm but not okay to select the better qualified candidate over the one who would have been better qualified had he not been unfairly disadvantaged by a systematically caused harm?

Second, this suspicion that the difference between systematic and non-systematic harm is irrelevant to the issue at hand seems to be vindicated if we again change the case of Bill and Ted to make it more like the case of the black and white candidates involved in affirmative action cases. The black candidates in question are hindered by harms that were inflicted systematically on many black families because they were black. So suppose that I change the story about Bill and Ted and make it the case that when Bill's father was framed, this wasn't a random act of wrongful prosecution, but was part of a systematic, large-scale enterprise. The government, let's suppose, had deliberately targeted Bill's father and many others like him, and it had organized the prosecutions because it didn't like something that all of the families had in common: they were all Mormons, let's say, or Irish. Does this change in the story change your intuition about the case?

In some respects, of course, the change in the story might well seem to make a moral difference. Robbing people because you want to get some money is wrong, for example, but robbing people because you want to harm them because of their religious or ethnic identity may well seem to be much worse. But that's not the relevant question here. The question is

not whether the change in the story changes your intuition about just how morally objectionable the act that caused the unfair disadvantage to Bill was. The question is whether the change in the story changes your intuition about what your moral obligations would be if you were in a position to select between the candidacies of Bill and Ted. If you're the track coach, for example, and if it's clear to you that Ted runs a little bit faster than Bill but that Bill would have been a bit faster than Ted if his father hadn't been unjustly prosecuted, do you think that it's okay for you to pick Ted over Bill without taking this fact into account if Bill's father was targeted randomly but not okay to do this if Bill's father was targeted because he's a Mormon? If you're the admissions officer and know that Ted has a somewhat better academic record than Bill but that Bill would have had a somewhat better academic record than Ted if his father had not been unjustly imprisoned, do you think that it's okay to simply admit Ted over Bill without taking this fact into account if Bill's father was a victim of a random injustice but not okay to do this if Bill's father was instead the victim of a plot against the Irish in particular? I, at least, find that these changes in the stories don't change my intuition about these questions. It seems just as clear to me in the revised version of the story as in my initial version that it's perfectly acceptable simply to select Ted over Bill even though Bill would have been better qualified than Ted had Bill's father not been unjustly imprisoned and even though Bill's father was targeted as part of a systematic wrong that was committed against members of Bill's ethnic or religious community. Again, it's hard to believe that many people on either side of the affirmative action debate would disagree. And so, again, it's hard to see how this second objection to my argument against the unfair disadvantage argument could be accepted.

A critic of my argument might also be bothered by a third difference between the two situations: in the case of Bill and Ted, the unfair disadvantage that holds Bill back a bit is the result of a recently committed injustice while the unfair disadvantage that holds many black candidates back a bit is the result of a series of injustices committed over a long period of time. But this worry, too, strikes me as ultimately unconvincing, and for the same two basic sorts of reason. First, it's once again hard to see why this difference should be morally relevant. In trying to figure out whether it's okay to select one particular candidate over another candidate who's a bit less qualified as a result of some systematically caused unfair disadvantage, why should it matter how long ago the harms that systematically caused the disadvantage began to occur? Again, it's not hard to see why the difference might matter if it had an effect on whether or not the disadvantage was caused fairly or unfairly, but if in both the case of Bill and Ted and the case of the somewhat less qualified black applicants and their somewhat better qualified white rivals, the systematically caused disadvantage is clearly an unfair one, why should it matter when the unfairness first began to occur? Why would it be okay simply to select the better qualified candidate over the one who would have been better qualified had he not been unfairly disadvantaged by a recent systematic injustice but not okay simply to select the better qualified candidate over the one who would have been better qualified had he not been unfairly disadvantaged by a longer-standing systematic injustice?

Second, this suspicion that the difference between recent systematic injustice and historical systematic injustice is irrelevant to the issue at hand seems to be vindicated if we once again change the case of Bill and Ted to make it more like the case of the black and white candidates involved in affirmative action cases. So suppose we start by going back a few centuries to the point at which a systematic injustice first starts to be aimed at one particular group of people. Let's say that a terrible feud breaks out between the Hatfields and the McCoys and that as a result, anyone who is a member of the McCoy family begins to be shunned and unfairly disadvantaged in other ways because of their biological identity as

members of this despised clan. A Hatfield and a McCoy are vying for the last spot in a college's freshman class and the Hatfield is a bit better qualified but only as result of the unfair disadvantage that the McCoy family has suffered.

If you've agreed with me up to this point, you'll presumably agree that it's morally permissible simply to admit the slightly better qualified Hatfield over the slightly less qualified McCoy without taking the consequences of the family feud into account. It might be nice or considerate of an admissions officer to take the McCoy applicant's family hardship into account, but morality doesn't require him to do so. But now let's say that twenty years or so have passed and that a second generation Hatfield and an unfairly disadvantaged second generation McCoy are competing for the last spot in a college's freshman class. Does the change from the first generation to the second generation change your moral intuition? To me, at least, the answer is clear: if it's morally acceptable to select the slightly better qualified first-generation Hatfield over the slightly less qualified and unfairly disadvantaged first-generation McCoy, then it's just as morally acceptable to select the slightly better qualified second-generation Hatfield over the slightly less qualified and unfairly disadvantaged second-generation McCoy. Once again, it's hard to believe that many people on either side of the affirmative action debate would think otherwise. And it's just as hard to see any morally relevant difference between the transition from the first generation to the second and the transition from the second generation to the third, the third to the fourth, and so on. As a result, it's hard to see any morally relevant difference between systematic unfair disadvantages that were caused recently and systematic unfair disadvantages that have been in place for a very long time. At least in terms of the question of whether we are obligated to give preferential consideration to those who are the victims of the unfair disadvantage, the difference in time seems to make no difference. And so this third objection, like the two before it, ultimately seems to be incapable of undermining my case against the unfair advantage argument.

Finally, a critic of my argument might point to the fact that the case of Bill and Ted involves an admissions officer making a decision about what to do in a particular instance while affirmative action involves a school making a decision about what sort of admissions policy to adopt and then apply to a large number of cases. And the critic might maintain that there is an important difference between the moral judgments we're inclined to make about particular cases involving particular people and the moral judgments we're entitled to make about what sorts of general policies should be adopted. Indeed, I suspect that a number of people who believe that affirmative action is morally obligatory will be wary of my attempt to argue from premises about particular cases to conclusions about general policies for precisely this reason.

But while I do appreciate this further difference between the two cases, it's once again difficult to see how the difference could be morally relevant. If it would be morally acceptable for a particular worker within a particular organization to behave in a particular way, after all, then why wouldn't it be morally acceptable for the organization itself to adopt a policy that would permit the worker to behave in that way? If it would be morally unacceptable for the worker to behave in a certain way, why wouldn't that make it morally unacceptable for the organization to adopt a policy that would permit it? The critic of my argument at this point concedes that it would be morally acceptable for you to simply admit Ted rather than Bill but denies that this judgment can be used to reach a conclusion about what sorts of affirmative action policies it would be morally acceptable for a school to adopt. But if it would be morally acceptable for you to admit Ted without adjusting for the fact that

Bill has been unfairly disadvantaged, then it would be morally acceptable for a school to adopt a policy that would allow its admissions officers to act in such a way in such cases. If it would be morally acceptable for a school to adopt such a policy, then it's morally acceptable for a school to adopt a policy that does not require its admissions officers to take unfair disadvantage into account. And if it's morally acceptable for a school to do this, then it's morally acceptable for a school to decline to practice race-based affirmative action.

When we set aside the divisive subject of race for a moment, I think that virtually everyone on both sides of the affirmative action debate agrees that we aren't morally obligated to favor a slightly less qualified white candidate over a slightly more qualified white candidate simply because the slightly less qualified white candidate has been disadvantaged in some way that isn't fair. And when we examine the matter more carefully, we see that virtually everyone on both sides of the affirmative action debate agrees about this even if there's a large number of such candidates rather than a small number, even if the cause of the unfair disadvantage has been systematic rather than non-systematic, even if the cause has been around for a long time rather than for a short time, and even if the question is about adopting a policy that permits such decisions to be made rather than about making such a decision in a particular instance. But given that all of this is true in the case where both candidates are white, all of this must be true in the case where one candidate is white and one is black. We therefore have good reason to reject the unfair disadvantage argument for race-based affirmative action on grounds that most people on both sides of the affirmative action debate already accept. The fact that black American candidates are unfairly disadvantaged by the legacy of slavery and its aftermath may make it desirable and even commendable for an organization to practice affirmative action. But it's not enough to make it morally obligatory.

objection two: the (other) compensation argument

A second kind of backwards-looking argument for an obligation to practice affirmative action appeals instead to a principle of compensation. The argument begins with a basic principle that most people are likely to accept. Suppose, for example, that Larry vandalizes Moe's car. It costs Moe \$3,000 to fix the damage and he has to spend \$250 renting a car while his is in the shop. Virtually everyone will agree that Larry owes Moe (at least) \$3,250 in compensation. This is because virtually everyone accepts the principle that when you wrongfully harm someone, you are morally obligated to compensate them for the harm. This means that you're obligated to restore your victim to the level of well-being that he would have enjoyed had you not wrongfully harmed him. But on the assumption that black Americans who apply for jobs and admission to schools have been disadvantaged by the legacy of slavery and its aftermath, this might well mean that such applicants are entitled to preferential treatment as a matter of compensatory justice. If a black job candidate would have been the best qualified for the job had he not been unjustly harmed, after all, then it would seem that giving him the job would be required in order to restore him to the level of well-being he would have enjoyed had he not been harmed in the first place. This appeal to compensation is perhaps less widespread than the appeal to fairness, but it provides an alternative argument for the obligation to practice affirmative action nonetheless, and one that at least some thinkers have begun to develop. In her provocative book *Racism and Justice: The Case For Affirmative Action*, for example, the philosopher Gertrude Ezorsky maintains that "blacks have a moral claim to compensation for past injury" and argues that affirmative action is thereby justified as "an appropriate method of compensation for blacks."¹⁵

¹⁵Ezorsky (1991: 73, 75). The same sort of argument is also defended in Masters (1990).

Like the unfair disadvantage argument, this compensation argument depends on two claims: the claim that black applicants for positions in the present are owed compensation for the harms that black Americans have suffered in the past, and the claim that if this is so, then they are entitled to preferential treatment when their applications are considered. As with the unfair disadvantage argument, I'm inclined to accept the first claim. Indeed, as with the unfair disadvantage argument, the first claim of this compensation argument is essentially the same as one of the claims that I depended on in my defense of slave reparations in chapters 2 and 3. But, again as with the unfair disadvantage argument, I'm inclined to reject the second claim that this compensation argument depends on. That's why although I eventually found myself convinced by the compensation argument for slave reparations, I continue to remain unconvinced by this other compensation argument for affirmative action.

To see what's wrong with the second claim needed to make the compensation argument for affirmative action work, it's important to begin by noticing an important difference between the unfair disadvantage argument and the compensation argument. The unfair disadvantage argument applies to every organization that accepts applications for some sort of position. That argument, after all, rests on the claim that if you hold a competition for a position, you're obligated to ensure that the results of the competition are fair. If the unfair disadvantage argument works, then, it will show that every organization should practice affirmative action. But the compensation argument, even if it turns out to be successful, can at most establish that an organization is obligated to engage in affirmative action if the organization itself can legitimately be held responsible for having caused the wrongful harm that rendered its black applicants a bit less qualified than they would otherwise be. If Larry vandalizes Moe's car, after all, then it's Larry that owes compensation to Moe, not Curly or Shemp. Similarly, if black candidates today are continuing to suffer the lingering consequences of slavery and its aftermath, then it is those who are responsible for causing slavery and its aftermath who would owe them compensation. Those schools or businesses that can't legitimately be held responsible for having caused these wrongful harms in the past can't be said to owe such compensation. And thus any argument for affirmative action that rests on a claim about such compensation can't apply to them.

With this limitation on the scope of the argument in mind, let's return to the case of Bill and Ted. But this time, let's say that both have submitted applications to be admitted to an educational institution run not privately, but by the federal government. If the compensation argument for affirmative action will work anywhere, it should at least work in this kind of case, since if any organization can be said to owe compensation to black Americans as a result of the legacy of legalized slavery, discrimination and segregation it's the federal government. So let's say that Bill and Ted have applied to West Point, and you're the head of the admissions office. Ted has a somewhat better academic and athletic record than Bill, but let's start by supposing that you know nothing about the circumstances that brought this about. Since Ted has the better record, you are about to offer him the final slot in the entering class. But at the last minute, you receive an additional piece of information: Bill's father was sent to prison a few years ago, and Bill's academic and athletic record was actually better than Ted's up to that point. It seems clear to you, in other words, that Bill's record would be a bit better than Ted's today if Bill's father had not been sent to prison.

As in the case of the unfair disadvantage argument, I think it's clear that most people, regardless of their views of race-based affirmative action, will agree that it's okay for you simply to admit Ted over Bill under these circumstances rather than to adjust your decision by taking Bill's misfortune into account. The fact that Bill would be mentally and physically

sharper than Ted today if his father hadn't been sent to prison doesn't mean that West Point is obligated to give him preference over Ted given that Ted is, in fact, mentally and physically sharper than Bill. But now let's add the fact that Bill's father was innocent and that he was unjustly framed and imprisoned by the federal government. Does this change in the story produce a change in your moral intuitions?

In some respects, of course, the change in the story does make a difference. You will now think that the government owes Bill's father compensation for the wrongful harm that it inflicted on him, and may well agree that it would also owe compensation to Bill himself for the harm that he suffered as a result of the harm that was inflicted on his father. But as the head of admissions at West Point, that's not the question that you have to answer. The question you have to answer is whether you should be obligated to admit Bill ahead of Ted. And with respect to this question, at least, it's hard to see how this change in the story could change your answer. If it's okay for you to admit Ted over Bill if Bill's father is guilty, why wouldn't it be okay to admit Ted over Bill if Bill's father is innocent? If you weren't obligated to admit Bill over Ted in the first place, then you aren't obligated to admit him now that you've discovered that it's the government's fault that his father was unjustly imprisoned. And since to me, at least, it seems clear that you weren't obligated to admit Bill in the first place, this means that you're still not obligated to admit Bill. It's true that Bill would have had somewhat higher grades and test scores and track records than Ted were it not for the terrible injustice committed by the government against his father. But this fact doesn't mean that you're morally obligated to admit Bill to West Point ahead of Ted.

Now let's assume that I'm right about that. If I am right, then an important lesson can be drawn from the case of Bill and Ted: although offers of admission to a state-run university (or of employment at state agencies) are a valuable benefit to those who receive them, the government is not obligated to distribute these benefits by giving them to the people who have been most harmed by the government. If Bill has suffered as a result of wrongful harms that the government inflicted on his father, then the government may well owe him some kind of compensation. But even if this is true, it doesn't mean that it's obligated to ensure that West Point admits him ahead of Ted. The compensation argument for affirmative action must therefore be rejected. A particular black student who is somewhat less qualified than a particular white student only because of the lingering effects of slavery and its aftermath, after all, is in precisely the same position that Bill was in with respect to Ted: he does not, in fact, have better grades and test scores than his rival but, by hypothesis, he would have had better grades and test scores were it not for the unjust harms that were committed by the government against members of previous generations of his family. Since this fact isn't enough to generate an obligation to admit the slightly less qualified Bill over the slightly more qualified Ted, it also isn't enough to generate an obligation to admit the slightly less qualified black student over the slightly more qualified white student.

My claim about this other compensation argument may at this point generate the same sort of confusion I tried to fend off near the beginning of this discussion. In chapters 2 and 3, I endorsed the claim that the federal government does have an obligation to make reparations to contemporary black Americans for the lingering harms caused by the legacy of slavery and its aftermath. But in this chapter, I'm denying that the federal government (or anyone else) has an obligation to benefit contemporary black Americans by engaging in race-based affirmative action. This may seem to be inconsistent. So let me offer one final example, once again trying to take race out of the picture for a moment. Suppose that the state of California wrongfully harms a white teenager by negligently causing him to be injured. The

white teenager sues for damages and a jury determines that he's entitled to a substantial benefit from the state. Suppose also that this teenager had applied to several campuses of the University of California school system. Based on his academic qualifications, he was rejected by Berkeley but admitted by a few second-tier schools. In this case, it should be clear that it would be consistent to maintain that the state should pay for the student's room and board at one of the schools he was admitted to and to deny that it should direct the Berkeley campus to admit him to their undergraduate program. But if that's right, then it's just as consistent for me to maintain what I have maintained here and in the previous chapters. A black student who gets into a second-tier campus but not into Berkeley might have a right to receive room and board support from the federal government as a form of reparations, but this would not mean that he has a right to have the government ensure that he be admitted to a school that he would not otherwise be admitted to. Most people seem to view the slave reparations position as an extreme, radical position while they view the claim that affirmative action is obligatory as an ordinary, common sense opinion. But once the arguments about compensation are carefully examined, we can see that just the opposite is true. The principle of compensation does provide a sound argument for the reparations position. But it doesn't provide a sound argument for the affirmative action position.

I conclude that backwards-looking arguments for affirmative action are unsuccessful. They may succeed in identifying some respect in which things would be better if a somewhat less qualified black candidate were selected over a somewhat more qualified white candidate. But even if they do, they fail to establish that there's a moral obligation to make things better in this respect. It's true, of course, that if a school declines to adopt an affirmative action policy and instead makes its admissions decisions in a race-neutral manner, then its policy will tend to work to the advantage of those who have already been unfairly advantaged and to the disadvantage of those who have already been disadvantaged. And it might be thought that this fact provides one final way to ground a backwards-looking defense of the claim that race-based affirmative action is morally obligatory. But to the extent that black Americans today are unfairly disadvantaged relative to white Americans on the whole, race-neutral policies which are uncontroversially just in all sorts of other contexts have this very same effect. Black Americans, for example, pay precisely the same tax rates that white Americans pay on income, property, purchases of goods, and so forth. They pay precisely the same fines for speeding, littering and any other fineable offense that white people pay. They are forced to comply with precisely the same burdensome licensing requirements that white people are forced to comply with in order to gain work as barbers, or doctors, or pilots, and so on. All of these race-neutral policies are better for white Americans and worse for black Americans than policies which took race into account by lowering the tax rates, fines, licensing requirements and so on for black Americans relative to those for white Americans. In all of these cases, then, the adoption of race-neutral policies perpetuates the currently existing inequities between black and white Americans. But in none of these contexts is this thought to make such race-neutral policies unjust. In the same way, then, and for the same reason, the fact that race-neutral admissions policies are better for white applicants and worse for black applicants than are race-based affirmative action policies doesn't mean that affirmative action is morally required or that doing without it is morally unjust. In the end, there seems to be no successful backwards-looking argument for the claim that race-based affirmative action is morally obligatory.

objection three: the appeal to diversity

Arguments for affirmative action that rest on considerations about compensation or fairness are essentially backwards-looking. They appeal to claims about things that happened in the past as a way of trying to justify an obligation we are said to have in the present. I've explained why I don't find these arguments convincing. But there's also a very different kind of argument that some people now make in defense of affirmative action. This more recent kind of argument is essentially forward-looking. It appeals to the claim that adopting affirmative action policies now will lead to valuable benefits in the future, and maintains that these benefits are sufficiently important to secure the case for affirmative action.

The most familiar version of this kind of argument involves an appeal to the value of diversity. Suppose, for example, that there are relatively few black students on a particular college campus. And suppose, also, that black students tend to come from backgrounds and have points of view that are distinct from those of typical white students. In that case, increasing the number of black students on the campus can reasonably be expected to help increase the extent to which students (white and black) are exposed to others from different backgrounds and with different points of view. Since being exposed to people with different backgrounds and perspectives seems likely to be of great educational value, this seems clearly to be the kind of thing that colleges should be doing. And so, according to this argument, it should seem pretty clear that affirmative action is the kind of thing that colleges should be doing as well. As Neil Rudenstine, the former president of Harvard, has put the point: diversity is "an educational resource comparable in importance to the faculty, library or science lab. . . . It is the substance from which much human learning, understanding, and wisdom derive. It offers one of the most powerful ways of creating the intellectual energy and robustness that lead to greater knowledge." University of Michigan President Lee Bollinger has argued that student diversity "is as essential as the study of the Middle Ages, of international politics and of Shakespeare" on the grounds that it "broadens the mind and the intellect – essential goals of education." And as Gerhard Casper, the former president of Stanford, put it a bit more bluntly, "We do not admit minorities to do them a favor."¹⁶

The appeal to diversity is perhaps the most common argument presented in defense of affirmative action on American campuses today. Not surprisingly, then, it's also one of the most controversial. Many people complain that the argument depends on the oversimplified (if not downright insulting) assumption that people of the same race all think alike. The philosopher Ellen Frankel Paul, for example, has objected that "the assumption that all blacks, all women, all Hispanics, and all members of other 'protected groups' share the same social philosophy is patronizing in the extreme and should be offensive to precisely those individuals it seeks to accommodate."¹⁷ This worry strikes me as overblown. The argument from diversity doesn't have to assume that all black people think alike and that all white people think alike. It simply has to assume that there are, on average, some differences in life experiences and points of view so that, on the whole, a group made up entirely of white Americans is likely to contain a narrower range of experiences and points of view than is a group that contains a representative portion of black Americans. And it isn't obvious, at least to me, that this is an unreasonable assumption.

¹⁶Rudenstine and Casper both quoted in Moore (2005: 8), Bollinger quoted in Arthur (2007: 263). For another useful defense of the diversity-based argument, see also Turnbull (1993: 139, 142).

¹⁷Paul (1993: 250). For similar complaints, see also Hartle (1993) and Wolf-Devine (1993: 230).

Other people maintain that while the argument from diversity might work in some contexts, it doesn't work in others. Maybe it's important to seek ethnic diversity in a history department, they might say, but it's not important in a math department.¹⁸ Different perspectives on the rise and fall of the Roman Empire are useful, after all. Different perspectives on the square root of seven are not. This objection, too, strikes me as unconvincing. Part of the reason that diversity might prove useful in a university setting does have to do with the importance of having different perspectives on a given subject. And it's probably (though not obviously) true that math and physics aren't subject to such different perspectives in the valuable way that history and literature are. But part of the reason that diversity might prove important has to do with its ability to foster innovative approaches to things other than the subject matter of a discipline itself: things like alternative approaches to teaching, to doing collaborative work with others, to resolving disputes with colleagues, to connecting the university to the larger community beyond, and so on. If, for example, black Americans on the whole are likely to have had different kinds of experiences interacting with authority figures than are white Americans, then they might, on average, be likely to come up with different approaches to understanding the authority that a teacher has over a student or an administrator over a teacher. There's no particular reason to think that a math teacher would benefit less from being exposed to more approaches to these concerns than would a history teacher. And so even if diversity might prove to be more valuable to a history department than to a math department, this strikes me as a poor reason to deny that it could be valuable to a math department at all.¹⁹

Understood as an argument for the claim that affirmative action has some positive consequences, then, I'm inclined to think that the appeal to diversity may well prove to be a reasonable one. And, indeed, its merits seem easy enough to apply to the world outside the university. In any kind of enterprise that benefits from innovation and the interplay of different ideas and approaches, some effort to create a more racially diverse population might prove capable of producing some tangible benefits.²⁰ But it's important to remember that the supporters of affirmative action that I'm concerned with in this chapter don't simply make the relatively modest claim that affirmative action has some positive consequences. Rather, they make the much more robust claim that there's a moral obligation to engage in affirmative action, that it would be positively immoral for a school or business or government to decline to have an affirmative action policy. And understood as an argument for the claim that there is such an obligation, the argument from diversity proves to be considerably weaker.

Perhaps the simplest way to see the problem with using the value of diversity as a basis for an obligation to engage in affirmative action is to look at other cases in which hiring someone would contribute to diversity. Suppose, for example, that a philosophy department is choosing between two finalists for a position teaching ethics. Both are white men and have

¹⁸See, e.g., Simon (1993).

¹⁹This reply also strikes me as a good reason to reject another objection that is sometimes raised against the diversity argument. Goldman, for example, complains that "proponents of [the argument from diversity] rely on correlations between race and gender and other factors such as diversity in academic viewpoints or expertise in such areas as feminism or African American studies. But when these other factors are easily measurable in themselves, as they are for potential faculty with written work to submit, then such statistical correlations should be irrelevant" (1993: 298). But this objection, too, unfairly limits the range of considerations that the diversity argument can appeal to. Whether or not a job candidate holds a certain view about the subject matter he teaches might be established by examining his written work, but whether or not his background and life experiences would lead him to approach his duties as a teacher, administrator or colleague differently cannot.

²⁰For a useful list of empirical work on the question of whether or not diversity produces such benefits in different areas, see Kellough (2006: 165n10).

very similar records, though one is just a little bit better qualified because he has published just a little bit more, or in slightly more impressive professional journals. The slightly better qualified applicant is from Chicago, and the slightly less qualified applicant was born, raised, and educated in Switzerland. The department currently has twelve members, all of whom have lived their entire lives in the United States. It seems clear in this case that the Swiss candidate would contribute an important and valuable kind of diversity to the department that the American candidate would not. People who were born and raised in other countries are, on average, likely to bring with them a distinctive set of life experiences and perspectives that can be beneficial to those who encounter them. And it seems plausible, at least to me, to think that this fact could count in favor of hiring him. But it's hard to imagine that anyone would think that this means the department has a moral *obligation* to hire the Swiss candidate over the American, that it would be positively unjust for it to hire the slightly better qualified American candidate instead.²¹

The same judgment, moreover, seems difficult to avoid in virtually every other context in which admitting or hiring someone would plausibly contribute to a valuable sort of diversity. Suppose, for example, that a department consists entirely of white male Christians and it is choosing between two white men, one of whom is also Christian, but the other of whom is Jewish or Muslim and a bit less qualified. Or suppose that a slightly less qualified candidate would bring diversity to a department by becoming the only member of the department to have been raised by a same-sex couple, or in an orphanage, or in extreme poverty, or to have served in the military, or in the Peace Corps, or in Congress, or to be deaf, or blind, or in a wheelchair, and so forth. In all of these cases, it seems plausible to me, at least, to think that the candidates who would add diversity really would bring something distinctive to the department in virtue of their unique background that would be of value to at least some of the students and faculty they would interact with. And so in all of these cases, it seems plausible, at least to me, to suppose that it could be appropriate to take diversity into account when making a final decision. But in none of these cases does it seem at all plausible to think that it would be positively unjust to refrain from doing so, and it's hard to believe that many people on either side of the race-based affirmative action debate would think otherwise. All of these cases, then, raise serious problems for the claim that the diversity that

²¹ Someone who agrees with this assessment might still want to endorse some kind of negative judgment here. They might say, for example, that the members of the department have a kind of professional responsibility to promote the intellectual development of the students and faculty and that if giving preference to the Swiss philosopher would better achieve this goal, then even if it isn't strictly speaking immoral or unjust to hire the slightly better qualified American philosopher, they still have a certain kind of responsibility to do so.

I have to admit that I'm a bit skeptical of this kind of argument. While I can see how it would be irresponsible for a department to allow its program to fall below a certain level of quality, it's less clear to me that it's irresponsible to decline to do everything one could do to maximize it. In this sense, there might be a professional responsibility to engage in affirmative action if the lack of diversity that would otherwise result would prevent the school from offering good quality programs, but no responsibility to do so if it could still offer perfectly good educational programs without doing so. And the burden would then be on the proponent of the argument not just to show that diversity can make things better, but that the absence of diversity prevents a school from meeting its professional duties to its students. It's much less clear that the argument could establish this.

But even if this kind of argument succeeds, it doesn't really undermine the main claim that I'm trying to advance in this chapter. It might help to show that the philosophers in a department that fails to promote diversity through affirmative action aren't doing their jobs very well, since they aren't fully living up to their professional responsibilities, but we don't generally think that it's immoral not to do your job well, and the main question I'm concerned with here is whether there's something immoral about not practicing affirmative action.

a black candidate would add to a department can by itself provide the basis for an obligation to give preference to that candidate.

This problem with the argument from diversity, in fact, runs even deeper. For there are still further sorts of cases where it's even more clear that hiring someone would increase diversity in a way that's valuable to a university, and yet even in these cases, it seems clear that this fact doesn't ground an obligation to promote diversity. Consider, for example, a department of East Asian Languages at a large university that's looking to hire another professor. The department already offers courses in Chinese, Japanese, Vietnamese and Thai, but offers no courses in Korean. The search produces two finalists: one who would offer some courses in Korean, and a slightly better qualified candidate who would offer courses in Chinese. In this case, it seems difficult to deny that the kind of diversity produced by hiring the Korean professor is a kind of curricular diversity that a university must place a high value on. Yet even here, it seems clear that this doesn't mean that the university has a moral obligation to hire the Korean teacher ahead of the Chinese one. If it decides to hire the slightly better qualified candidate rather than the one who would add more diversity to its curriculum, it doesn't act unjustly. It's difficult to imagine that any defenders of race-based affirmative action would disagree with this judgment. But if that's so, then even defenders of race-based affirmative action must concede that the claim that it's good to promote diversity isn't enough to justify the claim that there's a moral obligation to promote diversity. And if there's no such obligation, then while the appeal to diversity may well help to establish that affirmative action has some positive consequences, it will prove unable to establish that it's morally wrong to refrain from practicing affirmative action.

A defender of affirmative action, of course, could point to an obvious difference between promoting racial diversity and promoting the other sorts of diversity that the various examples that I have appealed to here involve. The lack of racial diversity in American institutions, after all, is often traceable to past acts of injustice in a way that the lack of these other forms of diversity typically isn't. That difference might well prove to be morally relevant. And so it could be perfectly consistent for a defender of race-based affirmative action to maintain that there is a moral obligation to promote racial diversity but no comparable moral obligation to promote, say, curricular or country of origin diversity.

But while appealing to this distinction might be relevant in attempting to defend race-based affirmative action by means of the kind of backwards-looking arguments that I considered in the previous two sections, it isn't available to the defender of affirmative action who seeks to provide an alternative kind of defense of the practice on forward-looking grounds. From that point of view, it doesn't matter whether an absence of diversity arose in a just or unjust manner in the past. All that matters is that an increase in diversity now will bring about benefits in the future. It's possible, of course, that people who defend affirmative action by appealing to the value of diversity would ultimately agree that promoting diversity is only obligatory when it's done in response to past injustices. And so it may be that seemingly forward-looking arguments for affirmative action ultimately collapse into the kind of backwards-looking arguments that I have already considered.²²

If the diversity argument is just another form of these other arguments, though, then it should be rejected for the reasons that I presented in the previous two sections. And if it isn't, then the diversity argument remains subject to the counterexamples I have employed in this section. And it will, in addition, have further implications that virtually every supporter

²²The philosopher George Sher, for example, argues that this is the case (1999).

of race-based affirmative action is likely to reject. If promoting diversity is obligatory even in cases where the current lack of diversity isn't a result of past injustice, for example, then historically black colleges like Morehouse and Spelman will have an obligation to increase their white enrolment even if this means admitting some white students ahead of some otherwise better qualified black students.²³ But it's difficult to imagine that many defenders of affirmative action would be willing to accept this result. Whether or not the argument from diversity amounts to another version of the argument from past injustice, then, it fails to provide a satisfactory justification for the claim that race-based affirmative action is morally obligatory.

objection four: the need for role models

Looking back at his lax study habits when he first arrived at college, Barack Obama has reflected on the important benefits he derived from having professors who were black like him:

I remember some of my teachers saying, 'Man, why are you pretending that you're not smart? Why are you spending all your time on something not serious, instead of focusing on what your talents are?' And, coming from black professors, especially, that was important, because I couldn't throw back at them, 'Oh, you just don't understand'. That's a big part of the reason it is so important to have black teachers, especially black male teachers. I'm not saying exclusively, but in many situations you need someone who can call you on your stuff and say . . . that it's not 'acting white' to read a book.²⁴

These considerations suggest a second and distinct way in which race-based affirmative action can be defended on forward-looking grounds. On this version of the argument, the important benefit that such policies can provide is a steady supply of role models who can inspire and motivate the younger black people who would look up to them. The philosopher Joel J. Kupperman, for example, has argued that disadvantaged minorities are often held back by things like fear of success and low expectations of what is possible. Seeing other people like themselves overcome such obstacles, he maintains, can help young black people follow suit. And so affirmative action programs are justified, on this account, by their ability to provide role models to those who really need them.²⁵

One possible problem with the role model argument, of course, is that it depends on the claim that role models really do provide the important benefits that defenders of the argument claim they do. Thomas Sowell, for example, has complained that there is no real empirical evidence to support this claim. If there are more black teachers at medical school, for example, will there eventually be more black doctors? Will there be better black doctors? It isn't easy to say.²⁶ And in a recent study on faculty diversity, Stephen Cole and Elinor Barber conclude that while having a professor as a role model can have a positive impact on a student's decision to pursue an academic career, the race of the faculty member is not relevant: "For students in all ethnic groups, it makes no difference whether the role model is of the same gender, the same race, or the same gender and race."²⁷

But even if the role model argument does succeed in proving that affirmative action produces some positive consequences, this isn't enough to show that affirmative action is

²³ I'm grateful to Jason Hanna for bringing this point to my attention.

²⁴ Quoted in Barkley (2005: 23-4). See also Thomas (1993).

²⁵ Kupperman (1993: 184).

²⁶ Sowell (2004: 144).

²⁷ Quoted by Arthur (2007: 261-2) from their book, *Increasing Faculty Diversity: The Occupational Choices of High-Achieving Minority Students* (Cambridge, MA: Harvard University Press, 2003), p. 247.

obligatory. This is the same basic feature that prevented the diversity argument from justifying an obligation to engage in affirmative action. There are, after all, lots of things that a university could do that would benefit a number of its students. It could build a bigger library. It could build a new recreation center. It could hire more people to work at its health center. It could offer better food in the dining halls. It could offer smaller classes. All of these would be good things to do because they would be good for the students. But the fact that they would be good things to do doesn't mean that it would be immoral not to do them. And so the same goes for the provision of role models for black students and others. The fact that it would be good to provide them isn't enough to make it immoral not to provide them.²⁸ And if the claim is that it's immoral not to provide the benefit of role models for students but not immoral not to provide the other benefits that I've appealed to here because the benefits of having role models are designed to respond to past injustices while the other benefits aren't, then while the claim might be relevant to the kind of backwards-looking argument that I considered in earlier sections, it won't be available to the defender of affirmative action who seeks to provide an alternative to those arguments.

objection five: the bias-elimination argument

Arguments for affirmative action that appeal to compensation or fairness are essentially backwards-looking. Arguments that appeal to diversity or the value of having role models are essentially forward-looking. I've explained why I don't think either of these kinds of arguments can provide a satisfactory justification for an obligation to engage in affirmative action. But there is one final direction in which a supporter of affirmative action might look: the present. There are, in particular, two kinds of argument that attempt to show that facts about the present suffice to generate an obligation to engage in some form of affirmative action. I will conclude my discussion of the arguments for affirmative action by considering them in turn.

The first kind of present-looking argument appeals to the claim that selection criteria that on the face of it may appear to be race-neutral are in fact biased against black people. If this is so, the argument maintains, then some form of affirmative action is required in order to compensate for this bias, regardless of its historical origins and regardless of its future consequences. Suppose, to take a familiar complaint, that the SAT is biased against black people.²⁹ Suppose, to be more specific, that the test is biased against them by 50 points. Perhaps the test includes vocabulary words that, on average, are more likely to be familiar to white students than to black students of comparable intelligence, or maybe it makes use of analogies involving examples that white students are more likely to be acquainted with than are comparably talented black students. A black student, on this account, is really just as

²⁸The philosopher Judith Jarvis Thomson makes this point using the example of the recreation center (1973: 368-9). The same problem ultimately undermines any attempt to ground an obligation to engage in affirmative action in the claim that such policies have positive consequences. On the account endorsed by Fullinwider and Lichtenberg (2004: 178, 208), for example, affirmative action is justified as a means of producing a greater amount of representation among black Americans as civic leaders, business leaders, educational leaders, and so on. But while this might again be a good state of affairs to bring about, the fact that it would be good cannot be enough to make it positively immoral not to bring it about.

²⁹Although this is a fairly common claim, it has long been accepted in the literature on academic testing that the SAT overpredicts rather than underpredicts the grades that black test takers go on to earn in the freshman year in college. On average, that is, the SAT actually overstates rather than understates the academic potential of the typical black test taker. Since colleges are primarily interested in the SAT to the extent that it provides a measure of how well a given student can reasonably be expected to perform in college, this suggests that, at least in terms of the way in which the test is actually used, the SAT is biased in favor of black students rather than against them (see Fullinwider and Lichtenberg 2004: 100).

smart as a white student who scores 50 points higher than him. And in that case, the argument maintains, a college admission policy should add 50 points to the test score of each black student before making a final admission decision. Affirmative action is justified in such cases simply as a means of assuring that the better candidate really does get admitted.³⁰

This particular version of the bias-elimination argument can apparently be strengthened by appealing to more general claims about biases against black Americans that currently exist. The social psychologist Claude Steele, for example, performed an experiment on two groups of students at Stanford University that seems to support the claim that what he calls “stereotype vulnerability” creates a general obstacle that black students confront, but white students do not, when being measured by apparently unbiased tests.³¹ Steele gave two groups of black and white students the same set of verbal questions that had been used in the GRE, the standardized SAT-like test that undergraduates who wish to attend graduate school take as part of the application process. Before taking the tests, one group was told that the exercise was simply intended to examine “psychological factors involved in solving verbal problems” while the other was told that it was “a genuine test of your verbal abilities and limitations”. In the first group, the black students scored as well as the white students. In the second group, the white students scored the same as the white students in the first group, but the scores of the black students were significantly lower. Steele takes this as evidence that when black students take tests such as the SAT, GRE and LSAT, they feel burdened by their attempt to overcome negative stereotypes about black academic ability. If this is correct, then even if the content of these tests is completely unbiased, the environment in which the tests are taken renders the results of the tests biased against black students. And if this is correct, then, again, it seems only fair to correct for the bias by giving black applicants some kind of compensating advantage in the way that race-based affirmative action programs seek to do.

Finally, some defenders of the bias-elimination argument for affirmative action appeal to the claim that the people who do the evaluating of applicants may themselves be biased against black candidates, even if only at an unconscious level. The philosopher Laurence Thomas, for example, defends the claim that an “unintentional” bias against black people is still widespread in the United States and that affirmative action is therefore justified as a means of “correcting for unfavorable moral head winds” of the sort that will otherwise unfairly bias a decision procedure in favor of white applicants over black applicants.³² Sometimes the evidence of such bias that is pointed to is simply the track record of the organization in question. If a law firm has a long record of consistently promoting its white associates to partner but of rarely if ever promoting its black associates, for example, this might be taken as reason to conclude that the people making the decisions harbor prejudices against black people, either consciously or unconsciously.³³ At other times, evidence of tacit bias against black candidates seems to emerge from more formal studies. In a 1985 experiment conducted by social psychologists Jeff Greenberg and Tom Pyszczynski, for

³⁰For some examples of the bias-elimination argument, see Francis (1993: 26-30) and Thomas (1993: 126-27).

³¹Armour (1997: 154-5) presents a useful summary of Steele’s research and appeals to it as support for the bias-elimination argument.

³² Thomas (1993: 126-27)

³³The philosopher James Rachels uses this kind of argument to defend the claim that a quota-based form of affirmative action might sometimes be justified (1993: 219-21). The “only purpose is to get the best qualified scholars” to fill the positions, he maintains, and “The question is simply what selection procedure will best serve that purpose. The fact of unconscious prejudice makes the usual system . . . an ineffective method” whereas imposing a quota “might be more effective because it reduces the influence of unconscious prejudice” (1993: 220).

example, different groups of white college students were asked to judge debates that were staged between black and white speakers. Immediately after the debate had ended, a person planted in the group by the experimenters did one of three things: he referred to the black speakers as “niggers”, criticized the black speakers in non-racial terms, or made no comment at all. Greenberg and Pyszczynki found that when judges overheard the racial insult, they had a significant tendency to lower their evaluation of the black speakers. They concluded from this that the judges did, in fact, harbor biases against the black speakers and that while the biases were not immediately obvious, the expression of racial slurs could, in effect, “cue prejudiced behavior in those who are exposed” to them and that this could have a negative impact on black people in such contexts as “parole board meetings, promotion committee meetings, and jury deliberation.”³⁴ Presumably, if these experiments do provide evidence of underlying biased tendencies, the biases would also be likely to have an effect on decisions regarding hiring and college admissions.

The bias-elimination argument depends on the claim that the deck is currently stacked against black candidates and that affirmative action is necessary in order to correct for this fact. One way to respond to the argument, then, would be to maintain that organizations can compensate for whatever bias that exists without engaging in affirmative action. The philosopher Lawrence C. Becker, for example, proposes that an organization could do without affirmative action by following this kind of policy: “Attend to the ways in which discriminatory bias may taint the use of formally neutral procedures and attend to the ways it may taint the definition of substantive criteria of competence. Correct for such bias...”³⁵ Similarly, Robert Simon has argued that “if we can make plausible judgments about when the qualifications of women and minority academics were underrated because of the illegitimate influence of stereotyping, then we may often be able to correct for such mistakes through such mechanisms as reviews and appeals.”³⁶

The problem with these sorts of suggestions, though, is that they assume that once we know that bias against black candidates exists, we can simply find it and take it into account. But part of the problem raised by the bias-elimination argument is precisely that the existence of a bias against black candidates can prevent us from fully recognizing its nature and reach. And so arguments of the sort made by Becker and Simon are unlikely to succeed on terms that proponents of the argument will be willing to accept. In order to show that there is a genuine problem with using the bias-elimination argument to justify an obligation to engage in affirmative action, then, it’s necessary to show that the argument is unsuccessful even if we concede that affirmative action is necessary in order to correct for bias that currently exists against black candidates. That’s what I want to do here.

Before presenting my principal objection to the bias-elimination argument, though, I want to briefly note one important preliminary concern: even if the bias-elimination argument proves to be successful, it isn’t completely clear that the practice it would justify would really be what most people on both sides of the debate mean by affirmative action. What people who oppose affirmative action typically oppose, for example, is the idea of selecting a black candidate over a white candidate when the white candidate is, all things considered, better qualified. They think, for example, that if a white student’s SAT scores are a bit higher than a black student’s, this is good evidence that the white student is a bit more academically

³⁴Quoted in Kennedy (2002: 60).

³⁵(1993: 98).

³⁶(1993: 73). See also Markie (1993: 281) for a similar suggestion.

skilled and so merits admission ahead of the black student for that reason. If the bias-elimination argument is successful, it will show that schools should in fact choose the black student over the white student in such cases. But since it will do so only by establishing that the black student really is, in fact, all things considered better qualified and more academically skilled, it isn't clear that the practice that it will justify is really the practice that opponents of affirmative action took themselves to be opposing. Similarly, what people who support affirmative action typically support is the idea that sometimes a black candidate should be selected over a white candidate even if the white candidate really is, all things considered, better qualified. They think, for example, that even if a black applicant is deemed to be a bit less academically skilled or qualified after all the relevant facts about such things as the accuracy of the SAT have been taken into account, he is entitled to be admitted ahead of a slightly more academically skilled or qualified white applicant because of the unfair circumstances that led him to be all things considered less skilled or qualified. If the bias-elimination argument is successful, it will show that schools should sometimes admit *seemingly* less qualified black applicants over *seemingly* more qualified white applicants. But since it will do so only in cases in which the seemingly less qualified black student proves, in fact, to be *actually* better qualified when all the relevant facts about such things as the accuracy of the SAT are taken into account, it isn't clear that the practice that it will justify is really the practice that supporters of affirmative action took themselves to be supporting. Even if the bias-elimination argument proves to be successful, then, it isn't clear how significant its success will be.

I'm inclined to think this is a significant problem for the bias-elimination argument. But a much bigger problem provides a reason to conclude that the argument is not, in fact, successful in generating an obligation to select the black candidate over the white candidate even in this more narrow range of cases. Suppose, for example, that Joe Black and Jack White are competing for the remaining spot in a particular graduate program. Suppose that Jack's GRE scores are a bit higher, that his letters of recommendation are all a bit stronger, and that the (all-white) group of professors from the department who meet with him find him a bit more impressive than they find Joe. But suppose also that Jack's modest edge over Joe in all of these respects is entirely due to subtle biases against black people that persist to this day: if the GRE were perfectly objective and administered in a perfectly race-neutral environment, then Joe would have outscored Jack. If their professors hadn't harbored vague and tacit stereotypical assumptions about black academic ability, they would have gotten to know Joe better during his college days and would have written more enthusiastic letters for him than for Jack. If the professors in the graduate program didn't share the same sorts of racial prejudices, they would have found Joe to be a bit more impressive than Jack when they met with him.

Given all of these suppositions, we can conclude that Joe Black is, in fact, a better candidate for admission than Jack White. We can also conclude that a variety of subtle biases against black people are preventing the graduate program's admissions committee from recognizing this fact. And we can even conclude that an affirmative action program that gave preference to a black candidate who seemed a bit less qualified over a white candidate who seemed a bit more qualified might well produce a better outcome as a result. We could agree, that is, that a graduate program that practices affirmative action under such conditions is more likely to end up admitting the better student than is a graduate program that doesn't practice affirmative action. Affirmative action, when practiced in a society that is biased against black people, could be understood as acting like a set of corrective lenses that simply help to produce the results that an admissions or hiring committee would be able to produce

without such a policy if its vision of things were not corrupted in various important respects. Affirmative action, in short, might be needed in order to make sure that the best candidate really does prevail.

Even if we were to accept all of this, however, this would still fail to establish that there is any moral obligation to practice affirmative action. The reason for this is simple: there is no moral obligation to select the best candidate in the first place. Thus, even if affirmative action really is needed in order to ensure that the best candidates are chosen, this can't show that there's anything immoral about a program or company or government that decides not to practice affirmative action. At most, it can show that there may be something unwise about their declining to engage in affirmative action. And that is a very different matter.³⁷

I'll try to defend my claim that there's no moral obligation to select the best candidate in just a moment. But before I do that, I think it's worth stressing that defenders of affirmative action, of all people, should be quick to embrace my claim anyhow. In most other contexts, after all, it's precisely the defenders of affirmative action who are prone to insist that it really would be okay to admit or hire a somewhat less qualified black person over a somewhat more qualified white person. But if they are right about this, then it clearly follows that there's no general moral obligation to always select the most qualified candidate. If there really were such an obligation, after all, then it would always be wrong to hire a less qualified black person over a more qualified white person. But while defenders of affirmative action seem clearly to recognize this point when they talk about the importance of compensation or fairness or about the value of diversity, they seem also to forget about it when something like the bias-elimination argument comes up. In the context of that argument, we're told that we have to engage in affirmative action to make sure that the better candidate wins. But if we don't really have to admit or hire the best candidate when the best candidate is white, then we also don't have to admit or hire the best candidate when the best candidate is black. Either way, if we want, morality allows us to choose to settle for less. And so, either way, morality won't require us to engage in affirmative action even if affirmative action really is needed in order to select the best overall candidates.

So I think that the defender of affirmative action, in particular, is already committed to the claim that I am using here to undermine the bias-elimination argument, the claim that there's no moral obligation to hire or admit the best qualified candidate. Since the claim that it's okay not to pick the best candidate will also play an important role in some of my responses to the arguments made by opponents of affirmative action a bit later in this chapter,

³⁷ A possible exception to this conclusion may be worth briefly noting: if a university has publicly committed itself to the claim that it only admits the very best students, then it may acquire an obligation to engage in a bias-correcting form of race-based affirmative action as a means of living up to its promise. A school that claimed to admit the smartest students but then knowingly ignored evidence that some of the black students it rejected were in fact smarter than some of the white students it accepted would therefore be guilty of acting in bad faith. While this kind of case is worth noting, however, its significance must remain limited in the context of the debate about affirmative action as a whole. The claim that a school could acquire an obligation to engage in affirmative action by acting in a certain way does not by itself provide any reason to think that there is an obligation to engage in affirmative action independent of such actions. After all, a school could clearly incur an obligation to engage in affirmative action simply by promising that it will engage in affirmative action, but the fact that it can be obligatory to do something if you promise to do it provides no reason to think that it is obligatory to do it even if you don't promise to do it. And the claim made by proponents of race-based affirmative action is the strong claim that schools and other organizations are obligated to engage in affirmative action whether they have already committed themselves to it or not.

though, I think it's important to see that the claim can be supported by appealing to other cases about which defenders and opponents of affirmative action alike will almost certainly agree.

Here's a simple case to start with. A tree blows over and makes a hole in the roof of my house. The hole is clearly visible from the street. Within a day, a number of roof repair companies have stopped by, taken a look, and left binding estimates for me explaining exactly what they would do to fix my roof, exactly how much they would charge me, exactly what their qualifications are, how many years of experience they have, testimonials from satisfied customers, and so forth. Now, am I morally obligated to look through all of their proposals and select the one that is the very best? Do I have to investigate the reputation of each company, the relative merits of the different materials they would use, and so on? Or would it be morally acceptable for me to look through the options, find one that seems good enough to me, and then just go with that one? Could I just pick one at random from the Yellow Pages? From a self-interested point of view, of course, it might well make sense for me to do a more thorough investigation. But morally speaking, it's very hard to believe that I'd be doing anything wrong if I simply picked the first offer that looked good enough to me and just got it over with, or even if I just selected someone at random. I think that virtually everyone, regardless of their views of affirmative action, would agree with this judgment. But this judgment clearly means that there's no moral obligation to select the most worthy person for the job.

A critic of my argument here might agree with my assessment of this particular case, but wonder whether it would hold in cases involving education or publicly funded organizations. So suppose, instead, that I am the chair of a philosophy department at a state university. The department suddenly finds that it is in need of an extra person to teach a logic class that begins in a few weeks, and I'm inundated with applications from over a hundred people. Am I morally obligated to carefully examine every detail in every application file so that I can pick the very best applicant, or is it okay if I simply offer the job to the first person I come across who is clearly qualified and clearly available? Again, it seems clear that morality doesn't require me to make sure that the very best candidate gets the job and so, again, it seems clear that even if affirmative action is necessary to make sure that the best candidate gets the job, this isn't enough to make affirmative action morally obligatory.

A final example is a bit more complicated, but worth considering because it's so widespread and so universally accepted. This is the practice that many colleges and universities have of making their admissions decisions on a rolling basis. When a school uses a standard admissions decision procedure, it operates on a deadline. If its applications in a given year are due on January 15, for example, then it waits until that date has passed, goes through all of the applications that it has received for that year, and picks out what it takes to be the best of all of them to make offers of admission to. When a school uses a rolling admissions policy instead, however, it has no such deadline. It evaluates each application as it comes in and makes a decision about whether or not the application seems good enough. As the year goes by, the number of available slots at the school goes down and so it typically becomes more and more difficult to get into that school. This means that a student who is admitted to the entering class of a school with a rolling admissions policy near the start of its admissions cycle may well have a weaker academic record than that of many of the students who are rejected for the very same entering class by the very same school near the end of the very same admissions cycle. Indeed, guidebooks for students applying to colleges frequently suggest that high school seniors with somewhat mediocre academic records deliberately seek

out schools that have rolling admissions policies and apply to them as early in the application cycle as possible for this very reason.

A school that adopts a rolling admissions policy does not, in fact, end up admitting the very best of the students who apply to it in a given year. If there's a moral obligation to admit only the best students who apply for admission to a given school in a given year, then rolling admissions policies are immoral. But while defenders and opponents of affirmative action may disagree about many things, I've never heard any of them maintain that rolling admissions policies are immoral. Nor does it seem to me at all plausible to think that they are. A rolling admissions policy enables a school to exercise more precise control over the exact number of students who end up enrolling in a given year than does a policy on which a school admits a group of students all at once and then sits back and waits to see how many people accept its offer. If a school wishes to put more weight on retaining this sort of predictability than on admitting only the very best students who apply in a given year, it seems clear to me that it is perfectly entitled to do so. This, again, means that it's morally okay to admit less than the best. And this, in turn, once again means that the bias-elimination argument is unsuccessful. There's no moral obligation to select the very best people who apply for any given position, and so even if existing biases against black people make it necessary to use affirmative action to accurately select the very best people, this fact provides no reason to believe that there's a moral obligation to engage in affirmative action.

A defender of affirmative action, of course, could once again appeal to the claim that the case of race is different because the history of racial injustices in this country makes it different. We have a moral obligation to correct for biases that arise from past injustices, on this account, but no moral obligation to correct for biases that arise in other ways. But while this kind of response would be relevant to those arguments that attempt to ground an obligation to engage in affirmative action in backwards-looking considerations about compensation or fairness, it isn't available to those who seek to provide an alternative to such arguments. If the fairness- or compensation- based arguments are successful, then there's no need to appeal to claims about currently-existing biases. And if those arguments are unsuccessful, then the fact that current racial biases are rooted in the past will prove to be irrelevant. If there's no obligation to take unfair disadvantage into account in the first place, for example, then the fact that currently-existing biases cause an unfair disadvantage to black candidates won't be relevant. And so, either way, the bias-elimination argument will fail to provide, or to support, a reason to conclude that race-based affirmative action is morally obligatory.

objection six: race as a qualification

A second kind of present-looking argument for affirmative action appeals to the claim that race itself can be a positive qualification for a position. Suppose, for example, that black students respond better to black teachers, or that black police officers can more effectively patrol predominantly black neighborhoods because they can elicit a higher degree of trust from the people who live there. If this is so, then it may well make sense to hire a black teacher or a black police officer with somewhat less experience or fewer credentials over an otherwise more qualified white applicant for the same position. Race-based affirmative action, in other words, might be called for in such cases to make sure that the candidate who

is hired is the one who really is the best qualified in the broadest and most important sense of the term.³⁸

I'm wary of the argument based on the claim that race is a qualification for several reasons. First, while it seems to be somewhat plausible in some cases, it seems clearly implausible in many others. Black students might well respond better to black teachers, for example, but it's very hard to see why, for example, black pilots or computer programmers or assembly-line workers would do their jobs any more effectively than their white counterparts simply by being black. Since most supporters of affirmative action for teachers and police officers are also supporters of affirmative action for such other positions as well, it seems hard to believe that the argument from race as a qualification can really serve as a satisfactory basis for their position.

A second and more worrisome problem with the view that race should be used as a qualification for a position is that it cuts both ways. Indeed, I'm inclined to think that, on the whole, putting this view into practice would do at least as much to harm black applicants as it would do to help them. Many schools in this country, for example, have virtually no black students. Imagine a black applicant for a job teaching kindergarten, or middle school, or college chemistry at such a school being told "I'm sorry, you were the smartest and best trained candidate we had, but we just don't have enough black students for you to teach here." Surely defenders of affirmative action would find this outrageous. But if race can be used as a qualification to justify hiring a black teacher because it is claimed that black students learn better from black teachers, then it's hard to see how we could avoid using race to justify hiring a white teacher on similar grounds in similar circumstances. And since there are far more white students in this country than there are black students, this kind of reasoning would tend to favor white teachers over black teachers more often than it would favor black teachers over white teachers.

The same kind of problem, moreover, seems to arise from the other sorts of cases that proponents of the argument from race as a qualification tend to give. If it's true that people who live in predominantly black neighborhoods are more likely to trust black police officers or black doctors, for example, then it seems at least as plausible to suppose that there are places where most of the people are white and are less comfortable around black people. It would then follow that white candidates should always be favored over black candidates for jobs that involve providing services in such areas. It's hard to believe that any advocate of affirmative action could accept this result. But if they're unwilling to accept it, then they should be unwilling to rely on the argument that race should be used as a qualification when trying to defend forms of affirmative action that would benefit black applicants. .

These two objections to the race as a qualification argument strike me as pretty significant. In addition, the argument seems to have the same basic problem that the bias-elimination argument has: even if it succeeds, it isn't at all clear that the practice it would justify is the one that defenders and opponents of affirmative action take themselves to be arguing about in the first place. Consider, for example, a case in which it seems clearly reasonable to treat sex, rather than race, as a qualification: a summer camp is hiring counselors who will, among other things, need to supervise young boys and girls while they're changing in and out of their swimsuits in the locker rooms at a public swimming pool. If there are going to be an equal number of boys and girls at the camp, it seems reasonable to

³⁸ This kind of argument is pressed by Shelton (1993: 235-38ff.).

ensure that there are a roughly equal number of male and female counselors. If most of the people who apply for those positions are male, then the camp may be forced to hire a female counselor with fewer years of experience and less glowing letters of reference over a more experienced and highly recommended male applicant. This is simply a case in which the logic of the race as qualification argument is applied to the case of sex. But hiring an otherwise less qualified female counselor because the camp needs another person to supervise the girls while they're in the locker room isn't the kind of thing that opponents of affirmative action take themselves to be opposing. And giving preferential treatment to female applicants only in cases where sex is clearly relevant in this way (the camp, let's say, gives no priority to female over male applicants when filling its administrative positions) isn't the kind of policy that defenders of affirmative action take themselves to be defending. Even if the race as qualification argument succeeds, then, it's not clear that it will justify the claim that supporters of race-based affirmative action are trying to justify or that it will be inconsistent with the claim that opponents of race-based affirmative action are trying to resist. In this sense, the race as qualification argument may prove in the end to be largely irrelevant.

But even if all of these concerns can be overcome, there's a much more fundamental problem with using the race as a qualification argument as a means of justifying an obligation to engage in race-based affirmative action. The problem is that, at most, the argument shows only that sometimes a black candidate who initially appears to be less qualified than a white candidate is, in fact, more qualified once all of the relevant considerations are taken into account. But even if this is true, this can be used as grounds for an obligation to hire the black candidate over the white candidate only if there's a general obligation to always hire the most qualified candidate. And as I pointed out in the previous section on the bias-elimination argument, it's implausible to suppose that there is any such obligation. There's nothing morally wrong with choosing to admit a less qualified student or to hire a less qualified worker. And so even if the race as a qualification argument could show that affirmative action is needed in order to identify the most qualified applicants, it still couldn't show that there's a moral obligation to engage in affirmative action.

I've now examined those justifications for affirmative action that focus on the past, the future, and the present. There is nowhere else to look. I conclude that while affirmative action may prove to have some desirable features, there's no reason to believe that it's morally obligatory. As far as morality and justice are concerned, if a school or business or government declines to practice affirmative action, that's okay.

objections from the right

I've argued so far that race-based affirmative action isn't morally obligatory. Many people will agree with me about this. But almost all of those who will agree with me will insist that I haven't gone nearly far enough. Not only is affirmative action not obligatory, they will say, it's not even permissible in the first place. It's not just that it isn't immoral not to practice affirmative action on their view. It's that it's positively immoral to practice it. But I think these people are mistaken, too. Having explained why I'm not convinced by arguments for the claim that affirmative action is obligatory, then, I'll now try to explain why I'm not convinced by arguments for the claim that it's impermissible. As with the case of the many arguments that have been offered in defense of affirmative action, it isn't always clear just

what, precisely, the many arguments that have been offered against affirmative action are trying to establish. Some arguments against affirmative action clearly aim to demonstrate that affirmative action is positively immoral. Others are somewhat ambiguous between claiming to show that affirmative action is immoral and merely claiming to show that it's a bad (but morally permissible) thing to do. And still others, perhaps, are offered as an attempt to identify something bad about affirmative action without really trying to show that it's wrong to practice it at all. Whatever the intentions behind the various arguments that have been made against affirmative action have been, though, they pose a challenge to my claim that affirmative action is permissible and non-obligatory only if they can do something to show that it's positively immoral to practice affirmative action. For the purposes of the discussion that follows, then, I'll focus on the question of whether any of them can.

objection seven: the right to be judged on one's individual merits

One kind of argument against affirmative action involves an appeal to people's rights. If affirmative action violates somebody's rights, then that's a good reason to conclude that the practice is unjust. And many people seem to believe that affirmative action does, in fact, violate somebody's rights. But while it's easy to find people who think that selecting a less qualified black applicant over a more qualified white applicant violates the white applicant's rights, it's often less easy to see exactly what right they're talking about. Early on in his book *Ending Affirmative Action: The Case for Colorblind Justice*, for example, the editor of the conservative *Weekly Standard* Terry Eastland claims that affirmative action policies always come "at the expense of individual rights."³⁹ But while the rest of his book spends plenty of time maintaining that affirmative action is a bad idea, at no point does it actually identify a specific individual right that white people have that's violated when they're passed over in favor of less qualified black people.

I'm inclined to think that there's a good reason for this omission. That's because I'm inclined to think that there's no such right. But since plenty of people seem to think that there is such a right, and since if there is such a right, it would certainly make for a powerful argument against the permissibility of affirmative action, I'll begin my survey of the arguments that people have made against affirmative action with this question: does affirmative action violate anyone's rights?

One common answer to this question is that affirmative action violates people's rights because people have the right to be judged on their individual merits. The argument for this position goes something like this: each person has the right to be judged as an individual, not merely as part of a group. When a person is judged as an individual, all that matters is that person's particular qualifications. So if each person is judged on their individual merits, the best qualified candidate will be selected. This means, in turn, that if each person truly does have a right to be judged on their individual merits, then the best qualified candidate really does have a right to be selected for the position he applies for. It follows from this that any decision procedure that deliberately selects a less qualified black applicant over a more qualified white applicant violates the better qualified white applicant's rights. And if that's true, then affirmative action is unjust and must be rejected for that reason.

Although this is one of the most common arguments against affirmative action, it's always struck me as one of the least plausible. This is because if the best qualified candidates really do have a right to be selected because their individual qualifications are the best, then

³⁹ Eastland (1997: 17)

all sorts of practices that virtually everyone accepts as perfectly permissible would turn out to be grossly immoral. Hardly anyone, for example, as I pointed out toward the beginning of this chapter, believes that geography-based affirmative action is unjust. If Yale or West Point admits a slightly less qualified candidate from North Dakota over a slightly more qualified candidate from New York, virtually no one believes that the New Yorker has had his rights violated. Even the famous Rhodes Scholarship program practices a form of geography-based affirmative action. Each year, it offers 32 scholarships to send students from the United States to Oxford University in England. But rather than picking the 32 best students who apply from across the country, it divides the country into regions and picks the best students from each region. A student from a highly competitive region who just misses winning a scholarship might well have won if he had instead come from a different and less competitive region and, again, no one really seems to believe that this practice violates anyone's rights. But if the New Yorker really does have the right to be judged solely on his individual qualifications, then it will turn out that his rights really are violated if he is passed over by Yale, or West Point, or the Rhodes Scholarship program for a slightly less qualified person from North Dakota. Since it's very hard to believe that he has had his rights violated in these cases, it's very hard to believe that he has a right to be judged solely on his individual merits in the first place.

In addition to the example of geography-based affirmative action, the various cases that I cited as problems for the bias-elimination argument a few sections ago pose problems for the argument based on a right to be judged solely on one's individual merits as well, and for the same reason. If I decide to hire someone to fix my roof who seems good enough, for example, rather than take the time to figure out which roofer is really the very best, virtually everyone would agree that I have done nothing morally wrong. If a college or university decides to admit students who seem good enough to it on a rolling admissions basis until all its slots are filled, rather than setting a deadline and then picking the very best of all the people who apply, virtually everyone will agree that it's acting within its rights. But if the very best candidate really does have the right to be selected, then I would violate the best roofer's rights by not selecting him and the school would violate the rights of those candidates who are rejected late in the admissions cycle when they are better qualified than some of those who were accepted earlier in the same cycle. Since these claims are extremely hard to believe, it's extremely hard to believe that you have the right to be selected if you are the best candidate. And if the best candidate does not, in fact, have a right to be selected simply because he's the best candidate, then the argument against affirmative action that's based on an appeal to the right to be judged on one's individual merits must be rejected.

I'm inclined to think that all of this is more than enough to warrant rejecting the argument against affirmative action that's based on the claim that people have a right to be judged exclusively on their individual merits. But just in case none of what I've said so far seems convincing, here's one final example. My neighborhood supermarket employs a number of people whose sole responsibility is to help bag the customer's groceries. The qualifications for this position seem pretty straightforward: you are good at this job to the extent that you can put things in the bags quickly, arrange them so that nothing breaks, do so without wasting too many bags, and be pleasant to the customers. Virtually all of the people who work there do just fine in all of these respects. One of the women who hold this position, though, doesn't do nearly as well. She is noticeably slower than the rest, often puts just one or two things in a bag before getting out another bag, is not as selective about which things go together in a bag, and is not particularly social with the customers, often failing to even make eye contact with them. This woman has Down syndrome.

Now, it's perfectly clear that this woman was hired ahead of other people who were better qualified for this particular position, but it's hard to imagine anyone thinking that the store acted immorally in hiring her because of this. If anything, I suspect that most people, regardless of their view of affirmative action, would admire the store for giving her a break despite the fact that she really isn't the most efficient bagger that it could hire. And yet, if the best qualified candidate really does have the right to be selected because people have a right to be judged exclusively on their individual merits, then this store, too, is guilty of acting unjustly. Since it will seem very clear to people on both sides of the affirmative action debate that the store hasn't acted unjustly in this case, it should seem very clear, once again, that the best qualified candidate for a position doesn't have the right to that position. And if that's so, then, once again, a white applicant who's passed over for a less qualified black applicant can't complain that his rights were violated simply because he was better qualified.

A critic of race-based affirmative action who agrees that it's permissible for the store to hire the inefficient bag checker with Down syndrome, of course, might try to account for this judgment by insisting that the woman really is the best qualified candidate overall. The critic might do this by arguing that efficiency is just one relevant qualification. If, for example, helping the store to maintain good relations with the local community is another qualification, and if occasionally hiring people with disabilities will help the store to do that, then perhaps the woman with Down syndrome really is the best qualified candidate after all and hiring her therefore doesn't violate the rights of any of the other candidates. But there are two problems with this attempt to fend off the case of the bag checker with Down syndrome. The first problem is that its justification of the store's right to hire her depends on the claim that the store in fact benefits from hiring her rather than a hiring more efficient bag checker. This means that the store wouldn't have the right to hire her if doing so would truly be altruistic, and it's hard to believe that many people on either side of the affirmative action debate would accept this result. The second problem is that if it's acceptable to insist that the woman with Down syndrome really is the best qualified candidate in this case, then it will be just as acceptable to insist that the seemingly less qualified black candidate really is the best qualified candidate in cases involving affirmative action. As long as the institution in question can point to some respect in which it thinks it will benefit by hiring or admitting the black candidate – increasing diversity, improving its relationship with the local black community, or feeling good about contributing to the rectification of past injustices for example – then it will turn out that the otherwise less qualified black applicant is in fact better qualified and that hiring or admitting him won't violate anyone else's rights after all. There seems, then, to be no credible way to reject affirmative action by appealing to the claim that the best qualified applicant has a right to the position.

objection eight: the right to equal treatment

A second and distinct right that critics of affirmative action frequently appeal to is the right to equal treatment. The argument based on this right is typically restricted to cases involving affirmative action as practiced by some branch of the government: affirmative action for police officers, postal workers, and so forth, or for students and faculty at state universities. So understood, the argument appeals to the claim that there is a right to "equal treatment" or "equal protection" under the law, and leaves open the possibility that it might be permissible for private companies and organizations to engage in affirmative action. The fiercely-debated Proposition 209, for example, which was passed in California in 1996, prohibited public schools like UCLA and Berkeley from practicing affirmative action in admissions and hiring, but didn't interfere with the right of private schools like Stanford and

Pomona to do so. And Ward Connerly, one of the most vocal leaders of the movement to endorse Proposition 209, declared in his victory speech on the night of the election that those who had voted for it had reaffirmed an inalienable “right” by proclaiming “the sanctity of the principle of equal treatment for all people.”⁴⁰ Following the position taken by those who supported Proposition 209 in California, and those who have supported similar movements in Texas and elsewhere, I will therefore limit my focus in this section to state-sponsored affirmative action. I want to consider the question of whether the appeal to a right to equal treatment can justify the claim that affirmative action is unjust at least in cases where it’s practiced by the government.

The answer to this question depends on what’s involved in having a right to equal treatment in the first place. So let’s suppose that you’re a white student who didn’t get into the local campus of the state university. And let’s suppose that you’ve since discovered that a number of black applicants were admitted to the campus during the same year you applied even though they had lower grades and test scores than you did. You’re wondering if the fact that the black applicants were given preferential treatment means that the state violated your right to equal treatment under the law. In order to answer this question, we first have to distinguish between a stronger and weaker sense in which you might be thought to have this right. And I think that once we draw this distinction, we’ll see that the claim that affirmative action violates the right to equal treatment is implausible.

Let’s start with a strong version of the right to equal treatment. On this understanding, it means the right to be treated in the same way as members of other races regardless of the circumstances. If you have a right to equal treatment in this strong sense, then it would always be wrong, and under any conditions, for the government to take your race into account in deciding how to treat you. Clearly a state university policy of giving preferential treatment to black applicants because they are black is inconsistent with this right. And so, if people really do have a right to equal treatment in this strong and unconditional sense, then that right is violated whenever a state university engages in affirmative action. And if such policies violate rights, they are unjust and must be eliminated for that reason.

But it’s implausible to believe that anyone has a right to equal treatment in this strong sense. Suppose, for example, that the CIA is considering two candidates for an undercover assignment in an overwhelmingly black African country. One of the candidates has a bit more experience and a bit more knowledge of surveillance techniques than the other, but the somewhat more experienced and knowledgeable agent is white while the somewhat less experienced and knowledgeable agent is black. Or suppose that a national park is auditioning two actors to portray Abraham Lincoln in an educational program for visiting school children. One of the candidates is a somewhat better actor with more performing experience, but he is black while the other candidate is white. Or the National Institutes of Health is looking to distribute some free medication as part of a study to determine whether or not a particular drug is especially effective in treating hypertension in African Americans. Virtually everyone, regardless of their view about race-based affirmative action, will agree that it’s morally permissible for the state to take race into account in these cases when deciding how to treat the people in question. But if the government has the right to take race into account in these cases, then there’s no right to equal treatment in the strong sense in which you have a right to be treated in the same way as members of other races regardless of

⁴⁰ (Connerly 1996: 65). Shelby Steele has also endorsed the claim that affirmative action violates the principle of equal protection under the law (Steele 2006: 35).

the circumstances. Any argument against affirmative action that requires the existence of the right in this extremely strong sense must therefore be rejected.

Those who support the equal treatment argument against affirmative action, of course, may not be particularly worried about these examples. “It’s fine to take race into account in those cases,” they might say, “because in those cases, the difference between black people and white people is relevant. When we said there was a right to equal treatment under the law, we didn’t mean that it was *always* wrong for the state to take race into account. We just meant that it was wrong for the state to take race into account when race is irrelevant. But race isn’t irrelevant in the cases you’ve presented. And so those cases don’t refute our claim that there’s a right to equal treatment.”

This seems like a perfectly reasonable response. On this account, the right to equal treatment doesn’t mean the right to be treated in the same way as members of other races regardless of the circumstances. Rather, it means the right to be treated in the same way as members of other races except in circumstances where differential treatment is justified. The claim that there’s a right to equal treatment in this weaker sense is therefore immune to these counterexamples. But now there is a new problem. In the case of the strong version of the right to equal treatment, it was clear that if there was such a right, then affirmative action would violate it. But in the case of the weak version of the right to equal treatment, this isn’t so. A policy violates someone’s right to equal treatment in the weak sense if it treats race as morally relevant under circumstances in which race isn’t morally relevant. And so affirmative action will violate the right to equal treatment in the weak sense only if it treats race as morally relevant when it isn’t.

But whether or not racial differences are morally relevant in the kinds of circumstances in which affirmative action programs exist is precisely the question at issue. Defenders of affirmative action in university admissions, for example, maintain that the difference between black applicants and white applicants is, in fact, morally relevant, perhaps because of differences in the obstacles they have had to overcome or in the life experiences that they can contribute to the campus environment. Opponents of affirmative action, on the other hand, maintain that such differences are morally irrelevant. But regardless of which side of the debate is right about the relevance of racial differences in these cases, the appeal to a right to equal treatment in the weak sense will prove unable to generate an argument against affirmative action. If defenders of affirmative action are correct when they claim that affirmative action appeals to considerations that are morally relevant, then affirmative action won’t violate the right to equal treatment in the weak sense. If opponents of affirmative action are correct when they claim that affirmative action fails to appeal to considerations that are morally relevant, then their case against affirmative action will already be complete and there will be no need to appeal to a right to equal treatment in the first place. Either way, the claim that there’s a right to equal treatment, understood in the weak sense, will do nothing to establish a case against affirmative action. If the right is strong enough to do the job of opposing affirmative action, it’s too strong to be plausible. If it’s weak enough to be plausible, it’s too weak to do any independent work.

These worries strike me as good enough reasons to reject the “equal treatment” argument against affirmative action. Before moving on to consider other arguments that have been raised against affirmative action, though, I’d like to briefly consider one response that might be made to the objection I’ve raised against this one, press one additional concern about the argument as I’ve presented it, and then say something about a somewhat different version of the argument.

The response that might be made to the objection I've raised so far is that there might be a still further interpretation of the right to equal treatment that's stronger than the weak version but weaker than the strong version. And, in particular, a defender of the right to equal treatment argument against affirmative action might well appeal to the doctrine of "strict scrutiny" that's been developed by the United States Supreme Court over the last few decades. According to this approach to understanding the rights guaranteed by the Fourteenth Amendment, the state can treat people differently according to race without violating the right to equal treatment as long as the state can establish that doing so "serves a compelling state interest" and does so in a way that's "necessary to serve that interest." It's plausible to think that this approach represents a reasonable middle ground between the weak and strong interpretations of the right that my objection appealed to, and plausible, too, to think that affirmative action fails to meet the test established by the strict scrutiny standard. If people have a right to equal treatment in this intermediate sense, and if affirmative action would violate the right in that sense, then the objection to affirmative action based on the existence of a right to equal treatment will prove to be successful after all.

The problem with this response to my objection, though, is that the strict scrutiny standard, too, has implications that virtually everyone on both sides of the affirmative action debate will reject. Consider, again, the case of the national park that wants to hire an actor to portray Abraham Lincoln in an educational program for school children. Or perhaps it wants to hire someone to portray Frederick Douglass, or to reenact scenes involving slaves and their masters. Virtually everyone would agree that it's morally permissible for the organizers of the event to take race into account when deciding which actors to cast in which roles. But it's clearly implausible to say that doing so is necessary for serving some compelling state interest. If people have a moral right not to have the state take their race into account unless doing so is necessary to serving some compelling state interest, then people have a moral right not to have race taken into account in this case. But virtually everyone, regardless of their view of racial profiling, will agree that people don't really have a moral right not to have race taken into account in this case. And so virtually everyone should agree that people don't really have the moral right to equal protection under the law in the way that the legal doctrine of strict scrutiny would interpret that right. And since people don't really have that right, the right can't be used to ground an argument against affirmative action.

The additional concern I want to raise about the equal treatment argument is this: if everyone really does have a right to equal treatment that can't be violated, and if taking race into account in hiring and admissions decisions really does turn out to violate this right, then why do people like Connerly think that private universities should still be allowed to engage in affirmative action? Consider, for example, a white student graduating from high school in California who applies to Stanford and to Berkeley. If Berkeley can't take his race into account because his right to equal treatment is simply inalienable, then why should Stanford be allowed to take his race into account? Either the right is inalienable or it isn't. If it's inalienable, then it seems that neither school should be allowed to violate it. If it isn't inalienable, then it seems that both schools should be allowed to violate it. Either way, it seems that if the kind of argument that people like Connerly have in mind is successful, then they should want the government to ban affirmative action in private schools (and businesses) as well. Some critics of affirmative action may accept this implication, of course, and so they won't see this as a further difficulty with the equal treatment argument.⁴¹ But most people

⁴¹ And Connerly, of course, may be one of them. Perhaps he thinks that affirmative action should be illegal at Stanford, too, and refrains from making this part of his position for reasons of political expediency. Still, the movement for Proposition 209 as a whole leaves the impression that affirmative action is okay in private schools

who make this kind of argument seem unwilling to have the government interfere with the marketplace by banning affirmative action in the private sphere. For those critics of affirmative action, at least, this provides one further reason to reject the equal treatment argument against affirmative action.

There is, however, a second way of understanding the argument based on a right to equal treatment. It seems to be less common in the academic literature on affirmative action, but my impression is that it's nonetheless quite popular in the public debate on the subject. On this second version of the argument, the reason that the difference between public and private institutions matters so much is that tax dollars go to support public institutions in a way that they don't go to support private institutions. Suppose, for example, that I'm a white high school senior in California and that I've applied both to Stanford and to the Berkeley campus of the University of California. Stanford isn't funded by my tax dollars and so, on this account, I don't have any particular right to have Stanford make its admissions decisions in one way rather than another. But Berkeley is supported by my tax dollars and so, the argument maintains, I have a right that I be given exactly the same amount of consideration as anyone else when Berkeley makes its decisions. I have to pay my taxes that go to Berkeley just like anyone else, and so Berkeley should have to treat me just like it would treat anyone else. If Berkeley gives preferential consideration to black applicants over white applicants, that is, then it violates my right to equal treatment in a way that Stanford's giving such preferential consideration wouldn't.

This second version of the right to equal treatment argument does seem able to avoid one potential difficulty with the first version, the difficulty of explaining why affirmative action should be banned at state but not at private universities. But in doing so, the second version of the argument creates an additional difficulty that the first version seems to avoid. The new problem arises when we compare cases involving students who apply to Berkeley from within California with cases involving students who apply to Berkeley from out of state. So suppose that Cal White and Cal Black are high school seniors in California and Kent White and Kent Black are high school seniors in Kentucky. Cal White is a bit better qualified than Cal Black and Kent White is a bit better qualified than Kent Black. All four apply to Berkeley.

This second version of the equal treatment argument seems to imply that while it would be wrong for Berkeley to admit the slightly less qualified Cal Black over the slightly more qualified Cal White, it wouldn't be wrong for Berkeley to admit the slightly less qualified Kent Black over the slightly more qualified Kent White. This is because the case that this version of the argument makes against affirmative action in public as opposed to private institutions is grounded in the claim that the white candidates who are disadvantaged by such forms of affirmative action have been paying taxes to support public but not private institutions. Cal White has been paying taxes in California and that's why he is supposed to have a right to equal consideration relative to Cal Black when Berkeley makes its admissions decisions. But Kent White and Kent Black haven't been paying any taxes in California, and so the argument provides no reason to think that Berkeley does something wrong if it admits the slightly less qualified Kent Black over the slightly more qualified Kent White. The second version of the equal treatment argument, in short, implies that state universities may not practice race-based affirmative action with respect to their in-state applicants, but may

but not in public schools, and this can still seem difficult to reconcile with the claim that affirmative action violates some inalienable right that people have to be treated equally by race.

practice race-based affirmative action with respect to their out-of-state applicants. I suspect that very few opponents of affirmative action would be willing to accept this result. Certainly the supporters of Proposition 209 didn't accept it. Proposition 209 abolished race-based affirmative action in state universities for all applicants, regardless of their state of origin. So for most opponents of affirmative action, the fact that the second version of the equal treatment argument has this implication should count as a good reason to reject it.

This first problem with the second version of the equal treatment argument strikes me as a pretty serious one. But there is, in any event, a second and much deeper problem with this version of the argument. The problem is that the argument neglects the distinction between having a right to equal treatment in two very different contexts: the context in which the government to which one pays taxes is deciding what kind of decision procedure one of its public institutions should use, and the context in which one of its public institutions is using such a decision procedure. As with many of the other problems that I've noted to this point, the importance of this distinction is best made clear by looking first at a case that doesn't involve race. And, again as with many of the other problems that I've noted to this point, once we look at such a case, I think it becomes clear that the argument in question has implications that are unacceptable to virtually everyone regardless of their views about race-based affirmative action.

Suppose, for example, that the government is trying to decide what sort of admissions policy the military academies should have. And suppose, in particular, that someone makes the following proposal: "people who live in the South aren't as important as people who live in the North, so we should have the military academies adopt the admissions policy that would best enable the military to serve the interests of people in the North." Clearly, in this case, the proposal would be unacceptable. The taxpayer in Alabama has a right to have his interests treated equally with those of the taxpayer in Maine when the government is deciding how to organize its military academies. But now suppose that everyone agrees that in deciding what sort of admission policy the military academies should have, the government must consider the interests of all Americans equally. And suppose that someone makes the following proposal: "since our military officers will eventually be forced to work and fight with people from all over the country, it's very important that they be exposed to people from all over the country during their training and education. Therefore, it's in the interest of all Americans, regardless of where in the county they live, that the military academies practice some form of geography-based affirmative action." This second proposal is clearly very different from the first. The first proposal involved justifying the adoption of a certain decision procedure by appealing to the claim that adopting the procedure would be good for people in one part of the country rather than in another. The second involves adopting a decision procedure that will itself give preference to applicants from some parts of the country rather than other parts, but it justifies adopting the procedure by showing that adopting it will be good for everyone, where each part of the country is valued equally. This is why virtually everyone would agree that the first proposal violates each American's right to equal treatment while the second proposal doesn't.

I think that virtually everyone on both sides of the race-based affirmative action debate will agree that the right to equal treatment that we have is a right that would protect us from the first kind of proposal but not the second. But if I'm right about that, then the second version of the equal treatment argument against race-based affirmative action should be rejected on grounds that virtually everyone on both sides of the debate already accepts. To see this, simply take the two proposals I mentioned above and change them to proposals

involving race-based affirmative action. Suppose, first, that someone made the following proposal: “black people are more important than white people, so Berkeley should choose the admissions policy that will best enable it to serve the interests of black people rather than white people.” Clearly this proposal would violate the right of white Californians to have their interests treated equally when deciding how the state should spend some of their tax money. But now suppose that everyone agrees that in deciding what admissions policy to have its public universities use, the state should treat the interests of black and white Californians equally. And suppose that someone then makes the following proposal: “since we want all of our students to learn as much as possible, and since on average a racially diverse campus will provide a wider variety of viewpoints and thus a more enriching learning environment than will a more racially homogenous campus, Berkeley should practice some form of race-based affirmative action.” If we agree that the second of the two proposals doesn’t violate the right to equal treatment in the case of geography-based affirmative action, then we must agree that the second of the two proposals here doesn’t violate the right to equal treatment in the case of race-based affirmative action.

If West Point admits a slightly less qualified candidate from North Dakota over a slightly better qualified candidate from New York, this doesn’t violate the New Yorker’s right to equal treatment despite the fact that some of his tax money goes to fund West Point. He has a right to have his interests considered equally when the decision is made about what kind of admissions policy West Point should have, but this doesn’t mean that West Point must have an admissions policy that treats him equally with the student from North Dakota. In just the same way, and for just the same reason, if Berkeley admits a slightly less qualified black Californian over a slightly more qualified white Californian, this doesn’t violate the white Californian’s right to equal treatment despite the fact that some of his tax money goes to fund Berkeley. He has the right to have his interests considered equally when the decision is made about what kind of admissions policy Berkeley should have, but this doesn’t mean that Berkeley must have an admissions policy that treats him equally with black Californians. In the end, then, neither version of the right to equal treatment argument provides a plausible reason to reject race-based affirmative action.

objection nine: the overgeneralizations argument

A related but distinct argument against affirmative action appeals to the claim that affirmative action depends on making overgeneralizations about people. Suppose, for example, that a state university admits some black students with grades and test scores that are lower than the grades and test scores of some of the white students that it rejects. And suppose, as is virtually always the case in such circumstances, that the school attempts to justify its right to do so by appealing to some claim like this: “black students have had to overcome race-based obstacles that white students have not had to overcome” or “black students would bring a distinctive point of view to campus that adding more white students would not”. In response to such claims, the argument against affirmative action that I want to focus on in this section complains that they rest on overgeneralizations. Yes, the critic of affirmative action will agree, many black students have had to overcome race-based obstacles, but not all of them have. And yes, they will continue, many black students would bring a distinctive point of view to campus, but not all of them would. And furthermore, they will add, some white students have had to overcome race-based obstacles and some would bring a distinctive point of view to campus. So while having been racially disadvantaged or having a distinctive point of view may in themselves prove to be relevant considerations, these critics will conclude, it would be wrong to give preference to black students over white

students on such grounds because this will provide a benefit to some students who don't really merit the benefit and will withhold a benefit from some students who really do merit the benefit. This argument is sometimes put in terms of the familiar language of rights ("I have the right to be judged as an individual, not as a member of a group") and sometimes put simply in terms of justice or fairness ("it isn't fair to lump me together with those other students just because of the color of my skin"). But no matter how, exactly, the argument is presented, it ultimately appeals to the claim that affirmative action rests on overgeneralizations and that this makes affirmative action wrong.

The problem with this overgeneralization argument, ironically, is that it depends on an overgeneralization -- an overgeneralization, moreover, about overgeneralizations. The argument, that is, depends on the claim that it's wrong to distinguish between two groups of people based on generalizations about the members of the groups. But while it's certainly wrong to do this sort of thing in some sorts of cases, there are plenty of other cases where it doesn't seem to be wrong at all. Consider, for example, distinctions that the law makes based on age. Most people below the age of eighteen fall below a certain level of maturity and most people eighteen and older meet that level or exceed it. So the state says that those who are not yet eighteen don't have the right to vote and those who are eighteen and older do. This rule clearly rests on a generalization: there are plenty of seventeen-year-olds who are mature enough to vote responsibly, and plenty of eighteen-year-olds who aren't. And the same is true of the many other laws that take age into account: laws regarding driving, smoking, drinking, marrying, engaging in sexual activity, serving in the military, and so forth. In all of these cases, the state draws a line between two groups based on generalizations about them. And, at least so far as I am aware, in all of these cases, opponents of affirmative action have never complained that the state's acting on these generalizations is impermissible. I've never heard a critic of affirmative action complain that the voting age violates the rights of mature seventeen-year-olds, for example, or argue that it's unfair for the law to allow immature fifty-year-olds to drink just because most fifty-year-olds are mature enough to drink responsibly.

But if it's permissible for the state to act on generalizations about age, why is it impermissible for it to act on generalizations about race? Indeed, if anything, the generalization that black students are more likely than white students to have suffered race-based obstacles in their upbringings is a far more reliable one than the generalization that teenagers one day past their eighteenth birthday are more likely to be mature enough to vote or smoke responsibly than are people one day before their eighteenth birthday. So far as acting on generalizations is concerned, then, it should be much harder to justify acting on age-based generalizations than to justify acting on the sorts of race-based generalizations that defenders of affirmative action typically appeal to. But since critics of affirmative action seem perfectly content to let the state act on generalizations based on age, they can't appeal to the fact that affirmative action also relies on generalizations as a reason for rejecting affirmative action.⁴²

⁴² As at least a few philosophers have pointed out, veteran's benefits provide another good example of this (see, e.g., Thomson (1973: 379-80), Ezorsky (1991: 75-6)). Virtually nobody complains if the state provides certain educational benefits to veterans that it does not provide to non-veterans, after all, since it seems fair to repay people who make significant sacrifices for their country. But not every veteran ends up having a bad or dangerous experience in the military, and some non-veterans make significant sacrifices for their country in other ways.

I'm inclined to think that this problem is enough to warrant dismissing the overgeneralizations argument. There is simply no general moral obligation to make sure that all and only those who deserve to be benefited are benefited, and so nothing particularly unjust about a policy that sometimes benefits those who don't deserve to be benefited and that sometimes overlooks those who do. But before moving on to other arguments against affirmative action, it's important to recognize that there's a way for the overgeneralizations argument to be reinforced. This involves appealing not simply to the claim that affirmative action fails to benefit some of those who deserve to be benefited, but to the stronger claim that it tends to benefit the least disadvantaged members of the groups that it does benefit and to do nothing for the most disadvantaged members of those groups. This is one of the central complaints about affirmative action made by some of its most prominent critics, including Thomas Sowell and Stephen Carter, and it's also one of the most common complaints about affirmative action contained in the literature as a whole.⁴³ So the overgeneralizations argument can't be fully dismissed without saying something about this further claim.

The first thing to say about this further claim is that it's not entirely clear that it's true. Proponents of this complaint typically point to cases involving affirmative action at elite universities and graduate schools. In order for affirmative action to help a black student get into Yale Law School, for example, the student must have a very strong undergraduate record. And, these critics point out, it's much more likely that a black student from a middle- or upper-class background will have attained such a record than that a black student raised in crushing poverty will have done so. Therefore, they conclude, affirmative action mostly benefits the best-off black candidates and does relatively little for the worst-off black candidates.

But this argument is unconvincing. It's true that if affirmative action is only practiced by elite schools, then it's likely to mostly provide benefits for students from relatively privileged backgrounds.⁴⁴ But that just seems like a reason to make sure that affirmative action is more widely practiced, not a reason to think that it's wrong to practice it in the first place. If schools that admit less advanced students to begin with and companies that hire workers with relatively few skills also practice affirmative action, that is, then it seems plausible to suppose that more black people from more seriously disadvantaged backgrounds will get diplomas and jobs as a result. When affirmative action policies were adopted by parts of the southern textile industry in the late 1960's, for example, the black workers who were hired generally did come from very disadvantaged backgrounds.⁴⁵ And there's no reason to think that this would not continue to be true today.

In addition, and perhaps more importantly, it's not at all obvious that affirmative action policies benefit only the particular black people who actually end up getting admitted or hired as a result of their being implemented. Affirmative action policies, after all, make things easier for every black applicant than it would otherwise be for that particular applicant,

⁴³ See, e.g., Sowell (2004: 12, 48, 76, 120, 186, 195) and Carter (1991: 71-2ff, 80, 133, 233). For other examples of this complaint about affirmative action, see also Kekes (1993: 150), Simon (1993: 57), and Kupperman (1993: 182, 183).

⁴⁴ Although it may be worth noting that when affirmative action enables a black student to enroll in an elite school that he might not otherwise have been admitted to, this means that he won't enroll at a less elite school that he would otherwise have ended up going to, and there will therefore be one more space at that less elite school for other black students to compete for, including students from less privileged backgrounds.

⁴⁵ Ezorsky cites this example as a response to the claim that affirmative action disproportionately benefits those who need it least (1991: 64).

and so they increase the chances of success for every black person. And, in general, it seems plausible to think that something that increases your chances of success is a benefit to you even if you don't end up succeeding. If using sunscreen increases your chances of successfully avoiding skin cancer, for example, then it seems clear that someone benefits you by giving you sunscreen to use even if you do end up getting skin cancer. Everyone who is given sunscreen is better off for getting it even if only some of them end up successfully avoiding skin cancer. In just the same way, it seems to me, every black person whose chances of success are increased by an affirmative action policy is better off because of the policy, even if only some of them end up successfully getting a job or a diploma as a result. So it seems to me that the popular claim so frequently used to prop up the overgeneralizations argument, the claim that affirmative action mostly benefits the least disadvantaged black people, is simply false.

But let's go ahead and assume that it's true. Let's suppose that the practice of affirmative action mostly benefits people who are already relatively well off. Even if this claim does turn out to be true, it's hard to see how it could count as a reason to think that affirmative action is immoral. It doesn't seem immoral in general, after all, to do things that primarily benefit people who are already doing pretty well. Consider, for example, an engineer who works for Rolls Royce, helping to make their cars run a little bit better. It seems plausible to suppose that most of the people who will benefit from what he does are already quite well off. But I doubt that anyone would conclude that what he does is therefore unjust. Or suppose that a federal law is passed that imposes stricter safety standards on the construction of luxury yachts. Again, it seems plausible to suppose that most of the people who would benefit from the law would already be quite well off. And, again, it seems unlikely that anyone, regardless of their view of affirmative action, would consider this a reason to deem the law unjust. The fact that doing something is likely to primarily benefit people who are already doing pretty well isn't a reason to think that doing it would be wrong. And so even if it's true that affirmative action primarily benefits people who are already doing pretty well, this isn't a reason to think that affirmative action is wrong.

objection ten: if discrimination was wrong then, then it's wrong now

The strongest arguments begin with the weakest assumptions. An assumption is weak if the claim it makes is so modest that virtually nobody would reject it. The weakness of an assumption makes an argument strong because it makes the argument's conclusion more difficult to resist. If you find yourself forced to accept an argument's starting point, after all, you may find yourself unable to avoid accepting its ending point, too.

Here's an example of a very weak assumption: the racial discrimination that white Americans imposed on black Americans in the past was morally wrong. Virtually everyone now agrees that this is true. Virtually no one today thinks that black people should be excluded from state universities, forced to drink from separate water fountains, or actively prevented from voting. If there were an argument that started from the assumption that these past forms of discrimination were wrong and that ended by showing that affirmative action today is wrong, it would be a very powerful argument indeed.

Many people seem to think that there is just such an argument. It goes like this: if it was wrong for racial discrimination to favor white people over black people in the past, then it's wrong for racial discrimination to favor black people over white people in the present. It was, in fact, wrong for racial discrimination to favor white people over black people in the past, so it is, in fact, wrong for racial discrimination to favor black people over white people

in the present. Affirmative action is a form of racial discrimination that favors black people over white people in the present. Therefore, affirmative action is wrong. As Supreme Court Justice Clarence Thomas has put it in advancing this kind of argument, “government-sponsored racial discrimination based on benign prejudice is just as noxious as discrimination inspired by malicious prejudice. In each instance, it is racial discrimination, plain and simple.”⁴⁶ As the philosopher Lisa Newton has put it, addressing affirmative action not just as practiced by the state but as practiced by schools and businesses, “When the southern employer refuses to hire blacks in white-collar jobs, when Wall Street will only hire women as secretaries with new titles, when Mississippi high schools routinely flunk all black boys above ninth grade, we have examples of injustice. . . . But, of course, when the employers and the schools *favor* women and blacks, the same injustice is done. Just as the previous discrimination did, this reverse discrimination violates the public equality which defines citizenship and destroys the rule of law for the areas in which these favors are granted.”⁴⁷ And as Virginia Black has put it even more forcefully, “If it is irrational and unjust and cruel to fire someone because he is a black or she is a woman – cases whose absurdity seems obvious – then it is equally irrational and unjust and cruel to hire someone because he is a black or she is a woman. To appreciate the parallel, one has only to remember that to hire X *because* of color is, ipso facto, *not* to hire Y because of color. When inscribed in law, this is racism.”⁴⁸

This is a very common argument against affirmative action. I don’t find it at all convincing. The problem with the argument is simple. It assumes that if an act is wrong in one kind of context, then it must be wrong in all other contexts. But that assumption is clearly false. Suppose, for example, that I asserted that a black person has the right to kill a white person in self-defense. If a white person has started shooting at a black person, for example, and if the only way for the black person to survive is to kill the white person who is attacking him, then the black person has the right to do so. I assume that virtually everyone would agree with this claim, regardless of their views of affirmative action. But now imagine that a critic of my assertion raised the following objection “if it’s okay for a black person to kill a white person in self-defense today, then it must have been okay for white people to kill black people by lynching them in the past. But everyone now agrees that it wasn’t okay for white people to kill black people by lynching them in the past. Therefore, your claim that it’s okay for a black person to kill a white person in self-defense in the present must be mistaken.”

In the case of this argument about the right to kill in self-defense, I think that virtually everyone would immediately see the problem with my critic’s position. My claim that it’s okay for a black person to kill a white person in one set of circumstances does commit me to something. It commits me to the claim that it would also be okay for a white person to kill a black person in the same circumstances. But just because it’s okay to kill in cases that involve self-defense doesn’t mean that it’s also okay to kill in cases that don’t involve self-defense. And so my critic would simply be wrong to maintain that my claim that black people may kill white people in self-defense in the present implies that it was okay for white people to kill black people by lynching them in the past.

⁴⁶ Quoted in Kennedy (1997: 6). Thomas has also written that “under our Constitution, the government may not make distinctions on the basis of race,” and so can also be understood as endorsing the strong (and, we saw earlier) implausible version of the right to equal protection argument (quoted in Fullinwider and Lichtenberg (2004: 179)).

⁴⁷ Newton (1973: 310).

⁴⁸ Black (1974: 106).

In the case I've just described, virtually everyone would see the flaw involved in the critic's reasoning. But many people nonetheless fail to recognize that precisely the same flaw is involved in the reasoning about affirmative action that appeals to the wrongness of past discrimination. When we all agree that the racial discrimination that favored white people over black people in the past was wrong, that is, this does commit us to something. It commits us to the view that it would also be wrong for racial discrimination to favor black people over white people in the same circumstances. Discrimination against black people in the past took place in a very specific context: it was used as a means of expressing the white majority's contempt for and feelings of superiority to the black minority. And so our rejection of this objectionable form of behavior does commit us to something: it commits us to the view that it would also be wrong for black people to use racial discrimination as a means of expressing contempt for or feelings of superiority to white people. But no one seriously believes that contemporary race-based affirmative action does anything of the sort. When the Governor of Alabama stood at the doors of his state's university and prevented black citizens of his own state from entering, his actions expressed the view that black people were incapable of, or unworthy of, higher education, that their presence would contaminate the environment in which their white counterparts were learning. This is what makes his act so deplorable. But if a state university today admits a somewhat less qualified black applicant over a somewhat more qualified white applicant, it's perfectly clear that the university's act in no way expresses contempt for or disapproval of the white applicant in particular or white people in general. And so the recognition that what was done by white people to black people in the past was seriously wrong provides no support for the conclusion that what is done to white people by affirmative action today is wrong.⁴⁹

I'm inclined to think that this is enough to show that the argument from the wrongness of past discrimination to the wrongness of present affirmative action is simply a bad argument. But since it seems to be such a popular argument against affirmative action, it may also be worth noting that the same reason for rejecting the argument can clearly be detected by looking at things from the other direction as well. The claim that discrimination that favors black people over white people today is okay, that is, does commit the defender of affirmative action to something. But what it commits him to is not the claim that discrimination that favors white people over black people would also be okay under any circumstances whatsoever. Rather, it commits him only to the claim that discrimination that favors white people over black people today would also be okay in the same sorts of circumstances. And this is an implication that surely any believer in affirmative action would find perfectly acceptable. If there were an African nation that had enslaved and persecuted a white minority for a substantial period of time, for example, and if the current generation of white people living there were significantly disadvantaged relative to the black population as a result, then a defender of affirmative action that favors black applicants over white applicants in the United States would have to agree that it would be permissible to practice a form of affirmative action that favored white applicants over black applicants in that country. If there's an academic discipline in this country in which white people are significantly underrepresented relative to black people, then, again, the defender of affirmative action would have to accept the claim that it would be permissible to favor white candidates over black candidates in such a case. If hardly any white students enter programs in black studies, for example, then the claim that affirmative action that favors black applicants over white applicants is appropriate in other departments might very well mean that affirmative action

⁴⁹ This point has been put nicely by the philosopher Thomas Nagel: "Traditional discrimination was as bad as it was not because it employed racial and sexual criteria, but because it told people they were despised or not taken seriously in light of their race or sex (1977: xii; see also 1973: 16).

that favors white applicants over black applicants would be appropriate in this department. But, again, it is hard to believe that any defenders of affirmative action would find these implications to be a problem. If supporting affirmative action for black Americans in the present meant supporting the wrongs that were committed against black Americans in the past, then there would certainly be a powerful argument against affirmative action lingering here. But it doesn't. So there isn't.

objection eleven: the bad track record argument

A related but distinct argument against affirmative action appeals to the claim that this country has a bad track record when it comes to treating people differently according to race. In the past, the argument points out, policies that treated people differently along racial lines almost always resulted in injustice. And because this is so, the argument concludes, policies shouldn't treat people differently along racial lines today. Affirmative action policies do treat people differently along racial lines today. And so, according to this argument, affirmative action policies are wrong. This seems, for example, to be the primary argument against affirmative action in Terry Eastland's lively book, *Ending Affirmative Action: The Case for Colorblind Justice*. In responding to the claim that racial classifications can be used by governments in permissible and even praiseworthy ways, Eastland characteristically writes that "we need only consider the history of our country to see that when government makes racial distinctions, bad things do follow."⁵⁰

The bad track record argument depends on the claim that if a distinction was frequently misused in the past, then it shouldn't be used today. But this claim is implausible. The fact that something used to be done incorrectly in the past provides no reason not to do it correctly today, though it may well provide reason to be cautious in attempting to do so. And the fact is that our government has a pretty bad track record with respect to all sorts of distinctions. The law, for example, distinguishes between adult and child, citizen and non-citizen, innocent and guilty, mentally competent and incompetent. In all of these cases, the distinctions have sometimes been misused in the past. But in none of these cases does it seem plausible to conclude that the government should therefore never draw the distinctions in the present. And so the same goes for the distinction between black and white. The fact that this distinction was misused in the past can't, by itself, give us a reason to think it must be wrong to use it differently in the present. It may give us reason to examine our thinking very carefully before we proceed, but if after examining all of the arguments about affirmative action soberly and dispassionately we can identify no other reason to believe that race-based affirmative action would be unjust, then the mere fact that in the past race-based distinctions were often misused shouldn't count as a reason to think that race-based affirmative action would be wrong.

objection twelve: the negatives of affirmative action

Many people who support affirmative action justify their support by appealing to the claim that affirmative action has positive consequences. As I argued earlier, this kind of approach may succeed in showing that affirmative action is a good idea, but it can't really show that an organization acts immorally if it decides not to engage in it. The fact that it would be good to do something, after all, is not by itself a reason to think that one has no right not to do it. Similarly, many people who oppose affirmative action justify their

⁵⁰ Eastland (1997: 141).

opposition by appealing to the claim that affirmative action has negative consequences. My argument in this section will be similar, too: while this kind of approach might succeed in showing that affirmative action is a bad idea, it can't really show that an organization acts immorally if it decides to engage in it. I will address each of the various negative consequences that critics of affirmative action appeal to one at a time. But the underlying problem with appealing to any of them is basically the same: the fact that an act will have some bad consequences isn't, in itself, a good reason to think that one has no right to do it.

Let's begin with perhaps the most common complaint of this sort: that affirmative action has a stigmatizing effect on the very people it's supposed to benefit.⁵¹ If everyone on a given college campus knows that the black students there were admitted under a relaxed set of standards relative to the white students, for example, then two things are likely to happen. First, the white students are likely to think less of the black students. Second, the black students are likely to think less of themselves. These strike many people as common sense claims, and there seems to be evidence available to support both of them.

In an experiment performed by Paul Sniderman and Thomas Piazza, for example, two groups of white people were asked whether they agreed or disagreed with three statements about black people: the claim that black people are irresponsible, the claim that they're lazy, and the claim that they're arrogant. One group was simply asked to respond to these three assertions. The other group was first told that a nearby state had set aside a "large number" of government jobs for black applicants "even if their scores on merit exams are lower than those of whites" and was asked to comment on this fact. The point of the experiment was not to compare the views held by supporters of affirmative action with the views held by their opponents. Rather, the point was to measure the effect that the mere mention of affirmative action, the mere awareness of its existence, might have on white attitudes toward black people regardless of their views about affirmative action. And the results were quite striking. When white people were simply asked to agree or disagree with the claim that black people are irresponsible, for example, 26 percent of them agreed. But when the existence of an affirmative action program was mentioned before they were asked the question, 43 percent agreed. When white people were simply asked to agree or disagree with the claim that black people are lazy, or arrogant, the rate of agreement was 20 percent and 29 percent respectively. When the affirmative action program was mentioned first, the numbers went up to 31 and 36 percent respectively.⁵² All of this seems to indicate that the existence of affirmative action programs makes white people less likely to think highly of black people.⁵³ Relatedly, there's no shortage of testimonies by black writers such as Stephen Carter, who writes of the "terrible psychological pressure that racial preferences often put on their beneficiaries."⁵⁴ These authors maintain that being admitted or employed as a black person under an affirmative action policy induces self-doubt. And since these people have first-hand experience with the problem they are describing, their testimony must be admitted to carry

⁵¹ See, e.g., Kekes (1993: 152-3), Paul (1993: 260), Eastland (1997: 9, 85-6, 141, 142), Steele (2006: 146; 2008: 13). For the related worry that it perpetuates the image of black people as victims, see also Pojman (1992: 188), and McWhorter (1996: 64ff.), which gives some evidence from the social sciences for this claim.

⁵² from *The Scar of Race* (get pp?) by Paul Sniderman and Thomas Piazza. McWhorter (1996: 71-2) appeals to one of their studies as evidence that affirmative action increases white hostility toward black people.

⁵³ Strictly speaking, I suppose, it might be argued that the studies only show that *mentioning* such programs has this effect. But it seems unlikely that mentioning the programs would have this effect if they did not exist.

⁵⁴ Carter (1991: 14) Though it is worth noting that Carter acknowledges that the self-doubts fade over time (1991: 24) and that Carter does not, in the end, oppose all forms of affirmative action (see, e.g., 1991: 84-5, where it is clear that he endorses at least a limited form of affirmative action in college admissions and perhaps even in some graduate programs).

great weight. On the whole, then, many people are likely to agree that affirmative action does, in fact, have these particular negative consequences.

I'm a philosopher, not a social scientist. So I'm not in a particularly privileged position to discuss the various pieces of empirical evidence that are typically presented in defense of the claim that affirmative action has a stigmatizing effect. But since philosophers are trained to critically analyze arguments that are built from such claims, I do feel competent to address the more fundamental question: even if we assume that affirmative action really is stigmatizing, does this provide a good reason to conclude that it's unjust? And with respect to this question, at least, it seems quite clear to me that the answer is no. Consider, for example, the practice that many universities have of admitting athletes under academic standards that are relaxed relative to the standards that are applied to non-athletes. Just as the average SAT score of the black students at a given school might be considerably lower than the average SAT score of the white students, that is, so the average SAT score of the white players on the school's football team might be considerably lower than the average SAT score of the rest of the white students on campus. In the case of reduced admissions standards for athletes, it again seems quite plausible to suppose that this will have a stigmatizing effect. Just as a white student who meets a black student may assume that the black student got in only because he's black, so a white non-athlete who meets a white athlete may assume that the white athlete got in only because he's an athlete. These forms of affirmative action, that is, may increase the number of people who think that black people are dumb and the number of people who think that athletes are dumb. Virtually nobody, though, really thinks that it's unjust for a university to give preference in admissions to athletes, even if doing so does have this stigmatizing effect, though of course many people think it's a bad idea. And the same is true for other commonly accepted forms of preferential treatment. Many universities, for example, give preference to so-called "legacies," the sons and daughters of their graduates. When I was a student at Yale, it wasn't uncommon for people to assume that such students had been admitted only because of who their parents were. And the same sort of thing can happen when schools practice geography-based forms of affirmative action. As all of these cases that don't involve race should make clear, the mere fact that an admissions policy may stigmatize the group it favors isn't enough to make it morally wrong. But if that's so when the favored groups aren't selected by race, then it's so when the favored groups are selected by race as well.

A second kind of negative consequence that critics of affirmative action typically appeal to concerns the effects that such programs may have on race relations. The argument here is that affirmative action is wrong because it is divisive, creating rifts on campus, in the workplace, and throughout society in general.⁵⁵ This kind of argument strikes me as puzzling. Affirmative action is certainly a divisive subject.⁵⁶ But it seems strange to blame one side rather than the other for this fact. The debate over abortion rights, for example, is at least as divisive as the debate over affirmative action. But imagine a defender of abortion rights suggesting that pro-life protesters should abandon their cause because the abortion controversy is tearing the nation apart. Surely, in that case, it would be clear that the pro-life

⁵⁵ See, e.g., Eastland (1997: 87), and Sowell, who argues that affirmative action has produced "Poisonous intergroup relations and real dangers to the fabric of society" in some countries Sowell (2004: 22). See also Sommers (1993: 292, 294).

⁵⁶ Note also that a 2000 *New York Times* poll asked "In order to make up for past discrimination, do you favor or oppose programs which make special efforts to help minorities to get ahead?" 46 percent of white respondents favored and 44 percent opposed, but for black respondents it was 76 percent in favor and only 11 percent opposed (Sack and Elder 2001: 372).

protesters would be entitled to reply by suggesting that the abortion rights advocates abandon their cause for precisely the same reason. It's true that if people stopped trying to outlaw abortion, the issue would no longer be so divisive. But it's equally true that if people stopped trying to prevent abortion from being outlawed, the issue would no longer be so divisive. The divisiveness of the abortion issue, in short, provides no reason to come down on one side of it rather than the other.

And yet while this fact seems to be clear to everyone in the context of the debate over abortion, many people seem to overlook it in the context of the debate over affirmative action. It's true that if supporters of affirmative action were to abandon their support, then many opponents of affirmative action would stop being so angry and resentful. But it's equally true that if opponents of affirmative action were to abandon their opposition, then many supporters of affirmative action would stop being so angry and resentful. Affirmative action programs make many people angry and resentful. Ending affirmative action programs makes many people angry and resentful. That's too bad. But it doesn't help to settle the controversy.

A third negative consequence that some critics of affirmative action appeal to is a reduction in efficiency. If a company fails to hire the best qualified candidate for any given position, then productivity will go down. If a college fails to admit the best students, then graduation rates – including black graduation rates -- will go down.⁵⁷ If a country's affirmative action policies make it too difficult for the best qualified white applicants to find employment, then they will look elsewhere. And so on.⁵⁸ Not surprisingly, defenders of affirmative action have produced a set of counterarguments to this claim that try to show that affirmative action actually increases efficiency. These arguments tend to appeal to the claim that affirmative action increases diversity and that diversity in turn tends to generate new ideas, new approaches, and so on.

I don't have much to say about whether or not affirmative action increases inefficiency. Sometimes it probably does, and sometimes it probably doesn't. But I do have something to say about whether or not increasing inefficiency would be enough to make affirmative action immoral. And what I have to say is this: it wouldn't. Indeed, once again, I think that this is already clear to most people on both sides of the affirmative action debate whenever they look at cases that don't involve race. When my neighborhood supermarket hired the grocery bagger with Down syndrome, for example, they failed to maximize the efficiency of their operation. But no one thinks that what the store did in hiring that worker was wrong for that reason because no one really thinks that there's a moral obligation to be as efficient as possible in the first place. And since it isn't wrong to fail to be as efficient as possible, the fact that affirmative action is inefficient, if it is a fact, isn't enough to make affirmative action wrong either.

Finally, it may be that some opponents of affirmative action consider the mere fact that such policies involve taking note of race itself to be a negative consequence of implementing them. Simply taking into account that one person is white and another person is black, after all, can be taken to reinforce the idea that there is indeed a difference between being white and being black and this, in turn, might be thought to impair our ability to work

⁵⁷ See, e.g., Arthur (2007: 277).

⁵⁸ The objection that affirmative action is inefficient has been pressed especially by Thomas Sowell in a number of works. See, e.g., Sowell (2004: 47-8)

toward a world where all people are treated the same. I'm a bit less clear about whether or not critics of affirmative action really mean to be pressing this claim, but whether there are people who mean to be asserting it or not, the claim itself certainly merits repudiation.

Virtually everyone, for example, agrees that it's permissible for the government to ensure that black people aren't excluded from serving on juries simply because they're black. In order to do so, the government must keep track of the number of white people and black people who live in a given area, the number of white people and black people who are called for jury duty, and the number of white people and black people who finally end up sitting on juries after they are called. In order to do all of this, the government clearly has to take note of who is white and who is black, as it must in any case where it attempts to enforce laws prohibiting racial discrimination in housing or employment. Virtually no one would say that the mere fact that the government must take note of who is white and who is black in these cases counts as a convincing argument against its doing so. And the same is true of the uncontroversial cases I appealed to in the section discussing the right to equal treatment argument. Virtually everyone on both sides of the affirmative action debate agrees that it's acceptable for the government to take race into account when looking for someone to go undercover in an overwhelmingly black country, to portray Abraham Lincoln in a play, or to take part in research on the effectiveness of certain drugs on African American patients. These actions, too, involve the government taking note of who is black and who is white and, again, since the consideration of race involved itself seems appropriate, the mere fact that they involve taking notice of race seems appropriate too. And so no one should be convinced by the claim that the mere fact that affirmative action involves taking note of who is white and who is black counts as an argument against it.

Like most policies, in any event, affirmative action programs have some positive consequences and some negative consequences. And, as with most policies, the positive consequences aren't enough to make affirmative action obligatory and the negative consequences aren't enough to make it unjust.

When a white employer has his son come to work for him at the family business, no one complains that he violates the rights of the more qualified candidates he might otherwise have hired. When the son's promotion within the company creates resentment among those who are passed over, when his co-workers assume that he has only gotten where he is because of his last name, or when the son himself begins to have doubts about his abilities because of the policy under which he was hired, no one concludes that these unfortunate consequences are enough to render the policy unjust. This is because, when we set aside the anger and emotion that racial issues often provoke in this country and focus on a case where a white employer hires a white applicant, we can see that it's simply okay to choose to hire someone who isn't the best candidate, even when doing so has these sorts of negative consequences. When a white employer hires a black applicant over other white candidates who are better qualified, however, many people lose sight of this important lesson. This is an unfortunate mistake. Ending affirmative action isn't wrong. But neither is not ending it.

a quick word about an irrelevant word

Before concluding this discussion, it may be worth briefly noting the fact that I've managed to write a whole chapter about affirmative action without once using the word "quota". How did I do that? The answer is simple: quotas are an irrelevant distraction. Although clearly

useful as a rhetorical device for rallying support among those who oppose affirmative action, the question of whether or not an affirmative action program employs numerical quotas makes no difference to the merit of any of the arguments that I have developed here.

That this is true can clearly be seen, as can so much else that is true in the debate over race-based affirmative action, by once again returning to the uncontroversial case of geography-based affirmative action. So suppose that two schools on the East Coast currently have no geography-based affirmative action program and suppose that both notice that they have very few students from North Dakota. Each decides that their overall academic mission would be better served if their students were exposed to people from a broader range of regional backgrounds, and so each decides to implement a geography-based affirmative action program. One school introduces an admissions policy that doesn't involve quotas: each qualified applicant from North Dakota has a few extra points added to the total score that the admissions office arrives at based on each student's grades and SAT scores. In some years, this results in two or three additional students being admitted from North Dakota. In some years, it results in seven or eight additional students being admitted. In some years, it turns out to make no difference at all. The other school introduces an admissions policy that does involve quotas. Since roughly two out of every thousand Americans live in North Dakota, and since the school admits roughly two thousand students each year, it sets a target of admitting four students from North Dakota each year: preference is given to qualified students from North Dakota over somewhat more qualified students from New York until four students from North Dakota have been admitted, at which point the rest of the applications from North Dakota are treated in the same way as the rest of the applications from elsewhere. In some years, the school doesn't receive enough qualified applications from North Dakota to fulfill its quota, in some years it receives so many qualified applications that it exceeds it, and in some years the applicants from North Dakota are so much weaker than the applicants from the New York that it ends up admitting no one from North Dakota at all.

In the case of these two schools and their geography-based affirmative action policies, two things should be clear. First, if the first school's policy is morally permissible, then so is the second school's. If a student from New York can't really complain that his rights are violated by the kind of geography-based affirmative action policy that doesn't involve a quota, that is, then there's no basis for complaining that his rights are violated by the kind of geography-based affirmative action policy that does involve a quota. And indeed, when schools like West Point distribute the majority of the nominations needed to secure admission to their program by congressional district in a way that clearly does amount to a quota system, no one really complains that this is unjust. As noted earlier, moreover, the Rhodes Scholarship program operates on a very precise quota system on which rather than selecting the 32 best candidates in the country, precisely four winners are selected from each of eight geographical districts. Again, virtually no one on either side of the affirmative action debate really thinks that this is unjust.

Second, if the first school's policy isn't morally obligatory, then neither is the second school's. If it would be morally permissible for the first school to abolish its non-quota program and revert to a policy on which it practices no geography-based affirmative action at all, that is, then it would be morally permissible for the second school to abolish its quota program and do the same. And, indeed, when schools refrain from practicing geography-based affirmative action, no one really complains that this is unjust, either. And surely no one would call it unjust if the Rhodes Scholarship program were to eliminate its geography-

based quota policy and were instead to simply select the 32 best candidates from across the country. If a geography-based affirmative action program that doesn't employ quotas is permissible but non-obligatory, in short, then a geography-based affirmative action program that does employ quotas is permissible but not-obligatory, too. Whether or not the policy in question makes use of numerical targets or quotas simply makes no moral difference.

But if this is true in the case of geography-based affirmative action, then there's no reason for it not to be true in the case of race-based affirmative action. None of the features that distinguish race-based affirmative action from geography-based affirmative action would suffice to render quotas impermissible in one case but permissible in the other or obligatory in one case but non-obligatory in the other. And so if race-based affirmative action without quotas is permissible but non-obligatory, as I have argued it is, then so is race-based affirmative action with quotas. Either way, race-based affirmative action is morally optional.

conclusions from the center

So what should we do? If either of the standard positions on affirmative action proved to be acceptable, this question would have an easy answer. If race-based affirmative action were unjust, then it would be wrong for any organization to practice it. If race-based affirmative action were obligatory, then it would be wrong for any organization not to practice it. But I don't think that either of the standard positions is acceptable. Race-based affirmative action, I've argued, is permissible but non-obligatory. It's morally optional.

To say that something is morally optional, however, isn't necessarily to say that it's morally neutral. Sometimes something is morally optional but still morally good. Donating some of your money to Oxfam, for example, might not be morally required of you, but even if it isn't, it's probably a praiseworthy thing for you to do nonetheless. Other things might be morally optional but still be morally bad. Morally speaking, for example, you may be within your rights to engage in true though mean-spirited gossip, but it might still be objectionable for you to exercise that right. And, of course, some things that are morally optional really do seem to be completely morally neutral as well. It's neither impermissible nor obligatory to put your left shoe on before your right, and it doesn't seem to be particularly good or bad to do so either. So if we accept my conclusion that race-based affirmative action is neither impermissible nor obligatory, we might still try to determine if it's morally good (though not required), morally bad (though not forbidden), or morally neutral (and thus essentially morally irrelevant). And in this way, perhaps, the defenders of the two more extreme positions that I've argued against in this chapter might still hold out some hope that, in the end, the truth will lie much closer to their own view than to the view of their opponents. Those who have sought to establish that affirmative action is obligatory, that is, might point out that what I've said in this chapter is consistent with the claim that it's morally much better to have affirmative action than not to have it, even if having it isn't strictly speaking required by morality. And those who have tried to show that affirmative action is unjust, on the other hand, might point out that what I've said in this chapter is consistent with the claim that it's morally much worse to have affirmative action than not to have it, even if having affirmative action is not strictly speaking forbidden by morality.

While either of these final judgments must be accepted as possibilities, though, they strike me as remote possibilities. As with most things that are morally optional, it seems unlikely that there will prove to be a simple, uniform answer to the question of whether it's just plain good or just plain bad to practice affirmative action. In this respect, demanding a simple yes or no answer to the question "should universities practice race-based affirmative

action?” is like demanding a general answer to the question “should universities have lacrosse teams?” In the case of the decision about having a lacrosse team, the requirements of morality will not provide an answer. Surely it’s morally permissible for a university to have a lacrosse team if it so chooses, and just as surely it’s morally permissible for a university to decline to have a lacrosse team if it so chooses. But given this, it also seems unlikely that there will be one generic answer to the question of what should be done that will suit all schools and all circumstances. In some cases, establishing a lacrosse team might well be a good thing to do, all things considered. It might benefit a group of students who are deserving young people and who might otherwise not get a chance to participate in an organized sport. It might expose others to a new and unfamiliar form of competition, with all the attending benefits that often come from increasing the diversity of one’s experiences. It might enhance the school’s profile. It might be fun. It might make people happy. In other cases, having a lacrosse team might prove to be a positively bad thing to do on the whole. The costs of running the program might be quite high. The marginal benefits to the participants might end up being relatively low. It might cause a good deal of resentment if it seems that funding has to be cut from some academic program in order to pay for the team’s uniforms. It might be boring. It might be dangerous. It might be divisive. It might make people unhappy. For these reasons, I doubt it would surprise anyone to conclude that there’s no general answer to the question of whether a university should have a lacrosse team, and that the only reasonable way of answering the question would be to examine the relevant details on a case by case basis.

But if what I’ve said about the morality of affirmative action in this chapter is correct, then we should say precisely the same thing about affirmative action: there’s no general answer to the question of whether an organization should engage in affirmative action, and the only reasonable way of answering the question would be to examine the relevant details on a case by case basis. In some cases, it might be good (but not obligatory) to have such policies. In other cases, it might be bad (but not unjust) to have them. Sometimes having such a policy may make no real difference at all and prove to be neither good nor bad.

So what should we do? We should resolve to accept that it’s always okay for an organization to practice affirmative action and that it’s always okay for an organization not to practice it. And we should direct our energies away from the often strident debate about whether affirmative action is obligatory or unjust and focus instead on the much more mundane, but much more important, task of trying to figure out which organizations would do more good than harm by having such a policy and which would do more harm than good.⁵⁹ I realize that this is a much less exciting answer to the question than the ones provided by the two dominant views on the subject of affirmative action. But sometimes, even for a philosopher, the less exciting answer to a question turns out to be the right one to accept.

⁵⁹ For a useful brief survey of some of the empirical work that has already been done on this question, see Kellough (2006: chapter 7).