

Two Approaches to Equality, with Implications for
Grutter

Boston University School of Law
Research Paper Series No. 22-39

November 3, 2022

Keith N. Hylton

Boston University School of Law

Two Approaches to Equality, with Implications for *Grutter*

Keith N. Hylton*

Abstract: The question “what is equality?”, applied to the distribution of resources across races, suggests the following answer: when there appears to be no need for a policy that focuses on improving the welfare of one race relative to another. There is another way to approach the same question: equality is when traditionally-recognized paths to advancement do not give preference to or disadvantage an individual because of his race. Notice the difference here is between *end-state* and *process-based* notions of equality, a distinction Nozick emphasized in his examination of justice in distribution. Nozick rejected end-state theories of justice in distribution. I side with Nozick’s approach and argue that the only morally justifiable and administratively feasible approach to determining equality in the distribution of resources across races is through a process-based definition. I explore the implications of this argument for *Grutter v. Bollinger*.

* Warren Distinguished Professor, Boston University, and Professor of Law, Boston University School of Law, knhylton@bu.edu. This paper was prepared from remarks made at the American Association of Law School’s Workshop on Racial Justice, Portland, Oregon. I thank David Lyons for comments on an earlier draft.

The question “what is equality?”, applied to the distribution of resources across races, suggests the following answer: when there appears to be no need for a policy that focuses on improving the welfare of one race relative to another. There is another way to approach the same question: equality is when traditionally recognized paths to advancement – what I will call *career channels* – do not give preference to or disadvantage an individual because of his race.

Notice the difference here is between *end-state* and *process-based* notions of equality, a distinction Nozick emphasized in his examination of justice in distribution.¹ Nozick rejected end-state theories of justice in distribution.² My inclination is to side with Nozick’s approach. I will argue below that the only morally justifiable and administratively feasible approach to determining equality in the distribution of resources across races is through a process-based definition. I will explore the implications of this argument for *Grutter v. Bollinger*,³ where the Supreme Court held that the University of Michigan’s use of race as a plus-factor in the law school admissions process was constitutionally permissible, and for the affirmative action debate generally.⁴ In the end, I find a paradox: the seemingly more conservative process-based notion of equality delivers a stronger defense of affirmative action than does the end-state approach.⁵

I will confine the discussion to two races: blacks and whites. However, the validity of the argument below does not depend on particular facts drawn from the history of relations between blacks and whites in the United States or anywhere else. The approach taken is equally applicable to questions about distributional justice across races in any setting.

Rejecting End-State Notions of Equality

¹ Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, Inc. 1974). Nozick’s book often refers to the distinction as one between “historical” and “end-state” (or end-result) principles. I prefer “process-based” in this setting, for reasons that will be clear in the text. The term “historical” suggests that the focus will be on groups and on the past, with suggestions for rectification of past injustices. The focus here is on the present and future.

² End state theories define inequality in terms of a particular desired distribution of a resource of resources. One might argue, for example, that the difference principle of Rawls is simply a special type of end-state theory of distributional justice. See John Rawls, *A Theory of Justice* (Cambridge, Mass: Harvard University Press, 1971).

³ 539 U.S. 306 (2003). Although *Grutter* was decided almost two decades ago, it remains of central importance in the affirmative action case law. In two new cases before the Supreme Court, *Students for Fair Admissions v. President & Fellows of Harvard College* and *Students for Fair Admissions v. University of North Carolina*, the plaintiffs are asking the Court to overturn *Grutter*.

⁴ The argument of this essay is in some respects similar to, but in other ways rejects, that of Dworkin. See Ronald Dworkin, *Taking Rights Seriously* 223-39 (1977) (chapter on affirmative action). The argument here is consistent with Dworkin’s argument that treatment as an equal is distinguishable and has different implications from the notion of equal treatment. However, I do not think it is possible to merge fairness concerns into a utilitarian analysis, as Dworkin does, in an adequately rigorous manner.

⁵ One could view this paradox as a special case of the more general point that the process-based approach to justice urged by Nozick does not necessarily imply a minimal state. See David Lyons, *Rights Against Humanity*, *The Philosophical Review*, vol.85 (2), 208-15 (April 1976).

What would equality in the end-state sense look like? One approach is to say that it is a state of “equal outcomes,” however that is defined. Another approach is to say that it is a state in which it appears that there are no advantages based on birth or status. The first approach is consistent with the end-state approach to determining equality. The second is closer to the process-based notion.

The equal-outcomes approach to determining inequality should be rejected, on administrative and moral grounds. To make this argument, I will set out a simple model of the equal-outcomes approach, in order to examine its necessary components. Although I take “equal outcomes” as the definition of equality in the end state, the argument in this section applies just as well to any other definition of the end state based on some particular distribution of outcomes.

First, the equal-outcomes approach requires an *equality monitor*. This person labels certain settings as “unequal,” others as “equal,” and, perhaps, presents a recommendation to the government to alter the outcomes in the unequal settings. To avoid any impression that I am biasing the argument against the equal-outcomes approach, I should add that an equality monitor is a general requirement of any scheme that attempts to determine whether resources or opportunities are distributed equally. In other words, to the extent that it is costly to have an equality monitor, it is an unavoidable cost. Any regime that attempts to define and address inequality will have to appoint an equality monitor.

Second, the equal-outcomes approach requires a definition of equality. The equality monitor would have to determine an end-state standard for measuring equality. For example, the monitor could examine the percentages of blacks and whites in certain positions, and label as unequal those cases in which the percentage of either group appears to be below some standard. The standard would have to be one that in the end-state, blacks do just as well as whites. The standard itself would require the monitor to choose a *base-line population*. The monitor would measure the percentages of blacks and whites in the base-line population and compare those to their respective percentages in the end state. If the percentage of blacks appeared to be below the standard adopted by the monitor, the end state would be labeled unequal.

Suppose blacks and whites invest in different ways in their careers. Suppose, for example, that 80 percent of whites go to medical school and 20 percent study art. Suppose among blacks, 40 percent go to medical school and 60 percent study art. In addition, assume the financial return from medical school is greater than that for art school. If the equality monitor demanded that blacks have the same income (or wealth) distribution as whites, he would have to transfer part of the return on the investment in medical school to blacks who went to art school. That would violate the expectations of whites who had invested in medical training, with anticipation of the usual reward, only to discover that they had received a lower return. It would also discourage whites from making investments into medical careers.

Now suppose the equality monitor says instead that blacks who go to medical school should have the same outcomes as whites who go to medical school, and blacks who go to art school should have the same outcomes as whites who go to art school. This means that if 20 percent of the whites who go to medical school make \$500,000 or more, the same percentage of blacks should have that income level.

This approach to determining equality is better than the first, but it still leads to transfers that would strike most to be unfair, in the sense of violating expectations, and could harm investment incentives. After all, not everyone invests the same amount in their medical school training, or in their careers. To impose equality on returns based on percentages that either entered or graduated from medical school would rob some medical school graduates of the return that usually accrues to extra effort invested during the years of training.

One counter to the unfairness argument is to say that it is too dependent upon individual expectations. If the government were to announce in advance that it would expropriate part of the reward from medical school to be redistributed on the basis of race, there would be no reason for white medical school graduates to feel that their expectations had been violated when the transfer occurred. If the moral case against expropriation under the “equal outcomes” approach depends on individual expectations, that case disappears, it would seem, when expectations are modified to incorporate the prospect of expropriation.

This counter to the unfairness argument, based on an expropriation preannouncement, must be rejected, for undermining the basis for property rights generally.⁶ First, the “expropriation announcement” would have to occur on some specific date, and on that date current medical school students would find their expectations violated, and their incentives to continue to invest in their careers correspondingly diminished. In other words, regardless of when the expropriation announcement is made, it would inevitably catch some medical students midstream in their investment period, violating their expectations. Second, for those individuals situated before the application stage to medical school, their expectations would be modified, as the preannouncement critique holds. But those individuals would learn from the example that another announcement could just as easily occur when they are midstream in the investment period, violating their modified (lessened) expectations. The reduction in the expected return, and the heightened risk of a further reduction, would discourage some individuals from going to medical school, and society would suffer to some degree from their discouragement.

⁶ The argument that individuals would adjust their expectations in response to an expropriation preannouncement reminds me of a series of arguments against property rights based loosely on the theory that people would readjust, and continue to invest, in the absence of rights, see Duncan Kennedy and Frank Michelman, *Are Property and Contract Efficient?*, 8 Hofstra Law Review 711 (1980). However, as an empirical matter, property rights support and enhance investment incentives. Douglass North and Robert Thomas, *The Rise of the Western World: A New Economic History*. New York: Cambridge University Press (1973).

Now, obviously, not all expectations present a moral case for being respected by the state. The expectation of a lobbying firm to get a monopoly, and the resulting windfall profit, through bribing state legislators does not present a moral case for respect – nor does the expectation of a thief with respect to stolen money. The moral case against expropriation or redistribution should not be viewed as resting entirely on expectations, without any regard to how they are formed. The moral case should be viewed as resting on *reasonable or legitimate expectations*, which incorporate an intuition for just rewards. Reasonable expectations of reward based on productive effort deserve to be respected by the state. Property rights, and quasi-property rights in the form of contractual obligations and market-based returns on investments in human capital, should be made secure in order to solidify the links between investment, expectation, and reward.

Of course, redistribution does not necessarily violate reasonable or legitimate expectations in every instance. For example, if redistribution corrects a failure in the market by internalizing an externality, it should not be viewed as violating legitimate expectations. Consider the redistribution that occurs when a tort victim is awarded a damage payment from an injurer. If the injurer acted negligently, he failed to properly take into account costs that he externalized to others while engaged in his activity. When a court awards a damage judgment, it redistributes resources from the injurer to the victim, but at the same time it internalizes an external cost generated by the injurer. Redistribution through internalization gives injurers incentives to take into account costs they externalize to others.⁷

However, redistribution or expropriation for its own sake, or to arbitrarily reassign a payoff from A to B, should be viewed as violating legitimate expectations. If the moral case against redistribution across races is viewed as resting on legitimate expectations, then merely introducing a redistribution program and announcing its existence should not be sufficient to destroy the moral objection to the expropriation of rewards.

Another counter to the fairness critique says that there is nothing wrong with the equal-outcomes approach, as long as we assume that there is discrimination in the real world. Equal outcomes aren't observed, one might argue, in part because whites have tilted the playing field to their favor. If whites have set up career-channel advantages that stay largely within their own race, then black medical school graduates will be denied opportunities to reach the same levels within the profession.

Introducing discrimination into the model immediately changes one's intuition on the morality of redistributing rewards on the basis of race, as would be required by the equal-outcomes approach. If some whites have gained rewards in part because of career-channel advantages designed to stay within their own race, it would not appear morally troubling to expropriate some of those rewards and redistribute them on the basis of race. But the equal-outcomes model does not provide such a surgical solution to the discrimination problem. It expropriates the rewards of whites whether or not they had

⁷ See generally, Guido Calabresi, *The Costs of Accidents: A Legal and Economic Analysis* (1970).

gained those rewards from discriminatory career-channel advantages and transfers them to blacks whether or not they had been victims of the same discriminatory processes. Because of this, introducing discrimination into the model does not provide a moral justification for redistribution on the basis of race.⁸ Moreover, the solution or “cure” to discrimination provided by the equal-outcomes approach may be ineffective. Demanding equal outcomes may entrench rather than lessen discrimination, and at the same time dampen investment incentives, for both blacks and whites.

In addition to all of this, there are administrative nightmares involved in the scenarios considered above. If the equality monitor seeks to avoid harming investment incentives, he will have to operate with the constraint that every individual gets the proper reward for his investments in human capital. How should one determine the proper reward for effort in training?

There is also the administrative problem of choosing a standard for determining inequality. Suppose the monitor defined an equal outcome as one in which black and white medical school graduates do equally well, where that means that the distribution of returns to a given investment path is the same across races. Thus, if 20 percent of white medical school graduates make over \$500,000 per year, 20 percent of black medical school graduates must make over \$500,000 per year for a state of equality to obtain. This approach to measuring equality involves both a choice of end state (making over \$500,000 per year) and base-line population (medical school graduates). Both choices are arbitrary.

Start with the choice of end state. If the monitor picked \$100,000 per year as his end-state measure, he might reach a different conclusion to his inequality inquiry than if he had chosen \$500,000. If so, how should he go about deciding which is the correct end-state measure? Suppose, for example, that for whites, the final income distribution is: 50% with \$50,000 per year, 50% with \$600,000 per year. Suppose for blacks the distribution is 10% with 50,000, 70% with \$200,000, and 20% with \$600,000. If the end-state standard is \$500,000, whites appear to be advantaged. If the end-state standard is \$100,000 blacks appear to be advantaged.

The choice of base-line population is still more troubling. In comparing outcomes of whites and blacks, should we compare the percentage of blacks who graduated from medical school to the percentage of white medical school graduates? Or, should we compare the percentages of white and black medical school graduates that *seek* jobs making over \$500,000 annually?

⁸ One might return to the initial description of the model and ask why blacks and whites would make different investments to begin with, if not for discrimination? Why would the percentage of whites choosing medical school over art school differ from the percentage of blacks making that choice? The model itself obviously provides no answer. However, if such a difference were observed in reality, I do not think that it would immediately imply some deeper level of discrimination. Preferences or cultural differences having nothing to do with discrimination might produce different investment decisions across groups. See generally, Thomas Sowell, *Markets and Minorities* (1981).

This is a bit like the nagging problem of market definition in antitrust law. The choice of market definition (base line) can determine the conclusion reached by the equality monitor. If he looks at the percentages that apply to high-paying jobs, and then selects a sub-group with certain skill qualifications, he may find no evidence of inequality. However, if he looks at the percentage that came out of medical school, the outcomes may look highly unequal. My suspicion, based in part on the antitrust analogy, is that there is no satisfying answer to this question.

Of the administrative issues facing the equality monitor, the base-line determination by itself has troubling moral implications. If, in comparison, we consider the other administrative problem, that of determining a fair return, we can at least say that there is a right answer to this question. It will be difficult to find the answer, and we may never be able to find with pinpoint accuracy the correct answer, but we do know that there is a right answer. We will have to live with an estimate that is hopefully within a tight statistical range of the right answer. The base-line problem is more troubling. At bottom, there may be no correct base-line population to use as a benchmark for determining inequality.

If there is no uniquely correct base-line population (nor even a small number of correct base-line measures) to use as a benchmark for determining inequality, then the work of the equality monitor who adopts the end-state approach builds on an arbitrary foundation. If the foundation is arbitrary, there is no moral basis for defending the equality monitor's conclusions. Moreover, the final result may be questionable because of the possibility of third-party influence. The equality monitor, after all, does not live in a vacuum. He lives in the real world, and has relationships, based on reciprocity, with others. If the foundation of his work is arbitrary, one can only wonder whether his recommendations are the result of a preference for some party who intervenes in his decision process.

To take a more concrete example, consider the case law professors are most familiar with, the law school appointments process. The equal outcomes approach would start with some statistical assessment of the shares of black and white applicants in the job pool for law teaching positions. Equal outcomes would require that the percentage of teaching positions should be the same as the shares in some base-line population. But what counts as a teaching position? Should we look only at the percentages getting tenure-track positions at elite schools? And what is the base-line population? Should it be determined by self-reporting, by looking at the people who file applications with the American Association of Law Schools? Or should the base line (job pool) be the pool of law school graduates with a certain grade-point average, whether or not they self-report as seekers of teaching jobs?

Now, I should say something about the implications of this argument. It is obviously an argument against quotas, since they tend to work on the same principle as the equal-outcomes approach. Less clear, the argument implies that even goals and timetables are difficult to defend. Any goal or timetable, to be effective, must be judged

against some numerical standard.⁹ But that takes us indirectly back to the equal-outcomes approach.

One possible implication is that a common method of determining discrimination – comparing percentages of certain ethnic or racial groups hired to their base-line populations in the job pool – may be difficult to defend. I am reluctant to draw this implication because the base-line comparison method often serves a different function in this setting. If someone is trying to determine whether discrimination has occurred, the base-line comparison approach is a useful starting point, from which a more detailed examination can begin. Moreover, in the discrimination context, the “end state” is not an arbitrary choice: the complainant identifies a selection process that he claims is discriminatory. Given that it has a more limited use and less arbitrary foundation, the base-line comparison approach is defensible as the starting point of an investigation into possible discrimination. However, if the base-line comparison approach is viewed as both the starting and ending point of such an investigation, then it clearly would be vulnerable to many of the objections set out above.

More generally, the use of a statistical measure as an *inference device* is distinguishable from its use as a way of defining or measuring inequality. In limited settings, statistical snap-shots based on comparing some end-state result with a base-line population can be helpful as a method of inferring discrimination or unfairness in a particular selection process. But even in these settings, the statistical measure merely gives rise to an inference of discrimination without providing proof. In the broader examination of distributional justice, end-state measures provide a considerably weaker basis for inference. Moreover, the statistical approach to finding discrimination boils down to a statistical proof that a given selection process operates unfairly – a proof that should be accomplished by statistically eliminating nondiscriminatory theories that might explain the results of the selection process. In other words, the statistical approach to determining discrimination has an ultimate goal that is equivalent to the process-based approach to determining unfairness.

Nozick’s argument against end-state definitions of justice focused on the instability of any particular distribution.¹⁰ Nozick noted that in a free market, people would make deals that violate the equal-outcomes standard, and there would be no evidence of anyone being worse off as a result. The same could be said about this case. Any equal outcomes approach ignores the effect of individual choice on the statistical picture that emerges. As a result, the ideal statistical picture becomes an elusive goal that requires constant intervention by the monitor to maintain.

Constant state intervention in order to maintain a statistical snap-shot of the desired end state involves the government in an effort to cancel market-determined outcomes. Let someone – say, a white doctor – develop a cure for cancer and make \$100 million as a result. That distorts the statistical picture, leading to a state of inequality.

⁹ For the argument that goals and timetables are equivalent to quotas, see Roland G. Fryer, Jr., 2009. *Implicit Quotas*, 38 *Journal of Legal Studies* 1 (2009).

¹⁰ Nozick, *supra* note 1, at 160-64.

The equality monitor would then have to intervene to “correct” the resulting distribution. How should it be corrected? There are many options available to the monitor, from leveling up, to the extent possible, to leveling down. All of them involve overturning outcomes based on voluntary contracts.

Process-Based

In contrast to the end-state notion of equality, there is a process-based notion that looks at whether lifetime channels for enhancement can be made equally available to all. Note that this would appear to many equality advocates to be the less effective approach toward achieving equality. The reason is that the process-based approach avoids any use of hard data on outcomes. An equality monitor under the process-based model does not ask for information on outcomes of blacks relative to whites. To many, this would seem to be a glaring example of a toothless program for achieving equality. However, I will argue that it does have teeth – that are likely to be more effective than those of the equal-outcomes approach.

So what would an equality monitor try to do under the process-based model? The answer depends on the setting. In general, the equality monitor takes an end state, and asks whether the processes leading to that end state provide equality of enhancement to blacks and whites. The equality monitor would not object to the processes merely because he observes different enhancement or access probabilities. Those different access probabilities might be based on investment decisions made along the way. The equality monitor would have to look for instances in which access probabilities differ in ways that have nothing to do with material investments made by candidates entering the process.

To lend some precision to this, let us define a career channel as consisting of periods of investment interrupted at points by *access nodes*. At each access node, the candidate is either accepted or rejected. If rejected, he continues to make investments toward some other access node. Over each period of investment, his likelihood of being accepted at the next access node is determined by his own investments and the investments of others. The equality monitor, under the process-based model, looks for instances in which access probabilities differ because of some exogenous obstacle at the access node itself (e.g., discrimination), or because of investments made by others (not the candidate) during the investment period.

Consider an example of one type of access node: a private firm’s hiring process. Is the hiring process one in which race would appear to be an irrelevant factor? How would we know if race is a relevant factor? There are certain features of a hiring process that would lend themselves to making race a relevant factor. To take an extreme example, suppose the firm showed a preference toward hiring friends and relatives of current employees. Since an individual’s friends and relatives are likely to be of the same race as the individual, a hiring process that favors friends and relatives of current employees makes race a relevant factor, and by doing so presents an obstacle to candidates of a different race from that of incumbent employees.

Another example: Abercrombie & Fitch, a clothes-maker that used to specialize in outdoor wear, but now specializes in casual fashion clothing for teens, made offers to customers to become salespeople and models for their store.¹¹ How did they go about choosing candidates? They looked for people who looked like those in their advertisements, which generally showed young whites in various states of undress. They tended to make their offers to whites who appeared to be physically attractive according to their preferences. The process led to very few, if any, offers to non-whites. Clearly, it was a hiring process in which race played a significant role. Whether it should concern the equality monitor is a different issue. An equality monitor might conclude that it is too trivial to be taken seriously as an issue, or that blacks who were cut out of the Abercrombie hiring process may have been better off as a result.

The examples I have just given – hiring friends and associates, hiring based on looks – are cases that would raise an eyebrow for the equality monitor, at the least. The monitor would probably feel a need to examine them closely to determine whether blacks and whites have equal chances of advancement. Suppose the monitor decides that in fact they do not have equal chances of advancement, that in each case the process is one that advantages whites. What should the monitor do next?

One answer is to do nothing, because the end result is trivial or unimportant. I gave the Abercrombie & Fitch example as one such case. If the monitor discovers that blacks had been passed over for the opportunity to be a salesperson or model for Abercrombie, he might not consider it a serious issue.

On the other hand, maybe it is a serious issue. Perhaps the equality monitor should be concerned with influencing the standards of attractiveness that would lead Abercrombie to implement a discriminatory hiring policy. But this approach would put the equality monitor onto an entirely new terrain, trying to alter tastes.

I have so far considered the altering of public preferences to be outside of the equality monitor's domain of concern. The Abercrombie example provides support for this assumption. It is a tall burden to change tastes, and the payoff in the Abercrombie case seems too small to justify the cost. However, there are other instances in which the altering of public preferences might have a substantial payoff. For example, the mutual fund industry depends heavily on getting people to believe that fund managers are competent managers of money. In general, investors in stocks do best by simply putting their money into an "unmanaged" index fund that is broadly diversified, like the S&P 500 index. The industry's welfare, however, depends on convincing ordinary investors that a typical money manager can beat the market, i.e., can perform better than a fund indexed to a broad basket of stocks such as the S&P 500. To convince investors that they should give their money to a fund manager, it helps to have managers who look like what an

¹¹ On the Abercrombie & Fitch discrimination story, see Steven Greenhouse, Clothing Chain Accused of Discrimination, N.Y. TIMES, June 17, 2003, at A1; TFL, How Abercrombie Ended up Being Sued by 250,000 Employees, <https://www.thefashionlaw.com/how-bercrombie-ended-up-being-sued-by-250000-employees/>.

uninformed investor would think a superior money manager would look like. This need to sell “image” provides a motivation to use race as a factor in the hiring process. The equality monitor may well conclude that this is a case in which it is worthwhile to intervene. Unlike the Abercrombie example, the mutual fund industry involves serious work that requires the development of useful skills. A hiring process that uses race as a factor provides differential incentives for the development of those skills, and supports a set of public preferences that have the same effect.

Having decided that a particular access node does not provide equal access probabilities to both blacks and whites, suppose the equality monitor decides that action should be taken to correct the unequal process. How should the process be corrected? The general goal should be to modify the process in such a way that the probabilities of advancement, summing over all of the access nodes, are the same for blacks and whites given the same individual investments.

The examples I have considered focus on access nodes – e.g., the job hiring process. The general focus here is on career channels, which consist of investment periods followed by access nodes. The equality monitor typically intervenes at an access node, and has to make a determination of whether the process is “equal” or “unequal”. Recall that the monitor finds the process unequal if he observes different probabilities of advancement at the access node that have nothing to do with the candidate’s own investments. If the equality monitor finds an unequal process and decides to intervene, he will do so with the goal of altering the career channels so that lifetime probabilities of enhancement are closer than in the original or undisturbed state.

The equality monitor has several options for intervention. In some cases, this will involve some policy that compensates for or mitigates an observed access differential. This option will change the “access process” so that it includes a mitigation policy. In other cases, the monitor may be able to reach back and influence the investment period itself to bring closer the acceptance probabilities at the upcoming access node. In terms of efficacy, it should be clear that it would be preferable to intervene, *ex ante*, at the investment stage rather than, *ex post*, at the access node. Intervening *ex ante* removes the need to intervene *ex post*, and its associated costs. However, *ex ante* intervention is prohibitively expensive in most cases.

Now, note that the first mitigation policy is one that some might call “affirmative action.” But the term has come to mean so much today that it does not always describe a policy of mitigating existing access differentials. Affirmative action suggests to many people the lowering of standards to allow the less competent to take the jobs of competent individuals. In other words, affirmative action has been understood to be the same thing as what an end-state equality monitor would impose.

But a mitigation policy is different from this notion of affirmative action. A mitigation policy is designed to modify a process that already yields different enhancement probabilities for blacks and whites, where those different probabilities are not based on the investments that the candidates made themselves. Thus, to alter the

process does not deny any candidate the reward that he reasonably anticipated as a result of his investments. It denies some candidates a gain that could be described as a windfall, and brings others to the level of return that is appropriate for their investment.

In general, a policy of mitigating access differentials is morally unobjectionable if it has the effect of wiping out windfalls and distributing them to deserving recipients. The same policy is objectionable if it destroys appropriately-anticipated returns for some and gives them as windfalls to others.

Application

There are two areas of life in which we see mitigation policies applied. One is the employment setting, the other the college application process. Most of the controversy over mitigation policies has involved college or graduate school applications, even though this is a small percentage of the cases in which these policies are applied. Mitigation policies are applied in the admissions processes of only a small number of colleges, perhaps the top 10 percent, affecting a tiny proportion of the population. Indeed, the controversy here is mostly a case of making an economic mountain out of an economic mole-hill. Still, this is what people want to argue about, so it is a worthwhile application of the model set out above.

The college and graduate school admissions process has become controversial because it is the one area in life where we see decision makers applying mitigating policies on a regular basis. As everyone knows, the top colleges use race as a plus-factor in the admissions process, along with other plus factors. I refer to this as a mitigating policy, but there are really two reasons that colleges use race as a factor in admissions. One is the mitigation theory, which is backward looking and external in effect. The other is a notion of diversity, which is forward looking and internal in effect.

The heavy use of mitigation in the admissions process distinguishes it from the employment context. Employers in competitive markets cannot afford to pay too much attention to mitigation. They will lose in the competitive market if they do. I would conjecture that the vast majority of mitigation that goes on in the private employment sector is insignificant. In many cases, diversity is the primary goal, which is almost always implemented as a rational, profit-maximizing policy. And studies show that “affirmative action” costs very little in the private employment sector.¹²

In addition to being a profit-maximizing policy in probably the vast majority of its uses, diversity is also distinguishable from mitigation because it is a utilitarian policy that has nothing to do with the question of justice in distribution. The diversity policy is typically adopted with some other payoff in mind. For example, an employer might

¹² Harry J. Holzer and David Neumark, *Assessing Affirmative Action*, 38 *Journal of Economic Literature* 483 (2000); Harry J. Holzer and David Neumark, *What Does Affirmative Action Do?* 53 *Journal of Industrial and Labor Relations Review* 240 (2000); Harry J. Holzer and David Neumark, *Are Affirmative Action Hires Less Qualified? Evidence from Employer-Employee Data on New Hires*, 17 *Journal of Labor Economics* 534 (1999).

adopt a diversity policy in order to communicate more effectively with its customers. A state agency might adopt the policy in order to provide better service to a diverse population of citizens. An educational institution might adopt the policy in order to provide what it views as a superior educational environment for its students.¹³ In each of these cases, the diversity policy is adopted not because diversity is an end in itself, but in order to achieve some other goal. Whether the diversity policy is socially desirable is an empirical question: does it accomplish the goal or goals for which the policy was adopted? If the policy fails to accomplish its goal, or accomplishes it at an unjustifiably large cost, then the policy fails on utilitarian grounds.

In the college admissions process mitigation means, in theory, applying a corrective policy to the process by which students gain a seat to a prestigious institution so that career channels for advancement appear to be closer to the equality standard. Needless to say, there are many problems thrown out by this statement. From which base line are access probabilities measured? In addition, how will the equality monitor avoid robbing some candidates of the return they rightly deserve?

In theory, access probabilities should consider the whole life of the candidate. There is no reason to measure access probabilities from, say, a month before the college application is due. For the same reason, four years before the application is due is an arbitrary cutoff. In other words, if, due to factors that have nothing to do with the candidate's own investments, he was at some point put on an inferior investment path than that of a rival candidate, this is a difference that the equality monitor should take into account.

Once the policy is stated in these general terms, there are many approaches that the equality monitor could take. If the equality monitor decides that blacks and whites do not have equal access he can try to design a mitigating policy that intervenes at any stage of the access process to bring about greater equality between the lifetime career channels.

How about the policy that we see enacted, of using race in certain cases as a plus-factor in college admissions? It is uncontroversially defensible under the theory set out. Given a differential in access probabilities, unrelated to individual investments, one is really looking at a distorted picture when the college applications come in. One imagines what a candidate from the relatively disadvantaged pool would have done if he had been on the same access path as one from the relatively advantaged pool. If the monitor concludes that he would have met or exceeded the acceptance criteria, then the equality monitor would enhance his probability of acceptance.

Now what does this imply for the affirmative action cases, the most important of which is *Grutter*, upholding the University of Michigan's use of race as a plus factor in the law school admissions process? It should be clear that this argument implies a highly individualistic approach to the admissions process, which I take to be the lesson of *Grutter*. Simply giving each candidate in the black pool 20 points, an approach the Court

¹³ *Grutter*, at 315.

invalidated in *Gratz v. Bollinger*,¹⁴ does not seek to determine whether each one presented a case that merited application of a mitigation policy. The core message of *Grutter* is entirely consistent with the process-based approach to ensuring equality.

The *Grutter* opinion also famously said that 25 years should be sufficient for a mitigation policy.¹⁵ This statement reflects some confusion on the Court's part between the process-based and end-state approaches. Under the end-state approach, a deadline would appear to be desirable. After all, the whole point of the equality monitor's intervention under the end-state model is to bring a desired statistical picture into reality. If the monitor is making progress, that statistical picture should be coming closer to reality, and so some sort of sunseting on the policy should be expected. If the monitor is making no progress at all toward the desired end state, the policy may be ineffective, which might make an even earlier termination date desirable.

Under the process model, there is no apparent need to put a termination date on a mitigation policy. The policy is applied on an individual basis. As long as the monitor finds evidence suggesting that factor X played a significant role in creating access differentials, the monitor is permitted to take factor X into account in his mitigation policy – whether factor X is race, poverty, geography, or whatever.¹⁶ Indeed, the notion of a time limit is inconsistent with *Grutter*'s core message of individualistic assessment. An individualistic assessment would remain precisely that no matter when it is carried out. To put a time limit on the policy undermines the justification for individualistic assessments.

Still, one might push back, even if a factor such as geography might play a role in creating access differentials, why should race play an independent role? It should not be controversial to note that there are ways in which race, alone, creates access differentials unrelated to geography or poverty. Race often induces a set of beliefs on the part of the perceiver and the perceived that affects interactions.¹⁷ Take the case of young school children. If a black student is enrolled in a poorly-performing urban public school, then race and geography could appear to combine so that it would be difficult to assign an independent effect to race. However, even in this case, race may have an independent

¹⁴ 539 U.S. 244 (2003).

¹⁵ *Grutter*, at 343.

¹⁶ One common counterargument is that geographic origin, or poverty, differs from race because geographic origin is not constitutionally prohibited as a factor of consideration in the admissions process while race is prohibited. This is false. Under existing law – consisting of *Regents of the University of California v. Bakke*, 438 U.S. 265 (1978); *Grutter*, and *Fisher v. Univ. of Tex. at Austin*, 579 U.S. 365 (2016) – race is not constitutionally prohibited as a factor of consideration. It is true that some people believe that the Constitution prohibits the use of race as a factor of consideration, but there are many who believe the opposite. Unless the Court were to attempt to operate as a legislature and count votes, such beliefs should play no role in an analysis of the legality of using race as a consideration factor.

¹⁷ See, for example, the literature on economics and identity. George A. Akerlof, Rachel E. Kranton, *Economics and Identity*, 115 *Quarterly Journal of Economics* 715 (2000); Richard H. McAdams, *Cooperation and Conflict: the Economics of Group Status Production and Race Discrimination*, 108 *Harvard Law Review* 1005 (1995); William A. Darity, Jr, Patrick L. Mason, and James B. Stewart, *The Economics of Identity: The Origin and Persistence of Racial Identity Norms*, 60 *Journal of Economic Behavior & Organization* 283 (2006).

effect,¹⁸ Suppose, instead, the black student is enrolled in a well-funded suburban school. Could race still operate to create access differentials? To take one example, if teachers in the suburban school differentially support and encourage students according to race, then clearly race could independently create access differentials even in the geographically-advantaged environment.¹⁹

Of course, in a zero-sum environment such as college admissions, adopting poverty, geography, or race as a plus factor implies a relative disadvantage to some other candidates. For example, awarding a preference for applicants from Appalachia implies a dispreference for candidates not from Appalachia. Such a dispreference would be morally objectionable if being raised in Appalachia failed generally to support an inference of an access differential. The question is whether “coming from Appalachia” is a sufficient statistic to support the inference of an access differential. By itself, it may not be: there is substantial income variation even in Appalachia. Indeed, none of the well-known plus factors is sufficient alone to infer the existence of an access differential. However, in combination with other data, such as on family income, “coming from Appalachia” may be sufficient to infer an access differential. Race is no different in this regard, and individualistic assessment processes of the sort suggested here would take into account race *and* factors other than race in a candidate’s assessment, just as they would take into account geographic origin *and* factors other than geographic origin in a candidate’s assessment. In such processes, taking factors such as Appalachian provenance and upbringing into account does not inappropriately transfer reasonably-expected rewards from some candidates and give them as windfalls to others. Indeed, it is the failure or refusal to consider Appalachian provenance and upbringing as a relevant statistic in the assessment process that would generate unjust rewards.²⁰

The *Grutter* Court’s finding that diversity provides a compelling justification for a policy that takes race into account is in palpable tension with the policy of individualistic assessment. The policy of individualistic assessment can be defended as a requirement of any serious effort toward justice in distribution. The diversity policy, unlike the process model, purchases its goal at the expense of distributional justice. If it has any impact at

¹⁸ See, for example, the “acting white” literature. E.g., Roland G. Fryer, Jr. and Paul Torelli, An empirical analysis of ‘acting white’, 94 *Journal of Public Economics* 380 (2010).

¹⁹ See Dorothy A. Cheng, *Teacher Racial Composition and Exclusion Rates among Black or African American Students*, 51 *Education and Urban Society* 822 (2019) (differential disciplinary exclusion rates and race); Seth Gershenson, Cassandra M. D. Hart, Joshua Hyman, Constance A. Lindsay, Nicholas W. Papageorge, *The Long-Run Impacts of Same-Race Teachers*, 14 *American Economic Journal: Economic Policy* 300 (2022); Thomas S. Dee, *A Teacher Like Me: Does Race, Ethnicity, or Gender Matter*, 95 *American Economic Review* 158 (2005) (differential educational encouragement and race).

²⁰ I recognize, of course, the truism that the law is not necessarily consonant with morality. See, e.g., Oliver Wendell Holmes, Jr., *The Path of the Law*, 10 *Harvard Law Review* 457 (1897). For example, the rescue, under reasonable conditions, of a person in danger of harm is the morally correct course of action. See, e.g., James Barr Ames, *Law and Morals*, 22 *Harvard Law Review* 97, 111-13 (1908). In spite of this, the law does not impose a duty to rescue: an excellent swimmer can watch a child drown in a pool and take no action, all without violating the law. *Osterlind v. Hill*, 160 N.E. 301 (Mass. 1928). However, the notion that the law *prohibits* the morally correct course of action, which the plaintiffs are in essence arguing before the Supreme Court in *Students for Fair Admissions v. President & Fellows of Harvard College* and *Students for Fair Admissions v. University of North Carolina*, is a different concept altogether.

all on the selection process beyond that of the mitigation policy, the diversity policy involves sacrificing distributional justice or fairness in order to move closer to a statistical snap-shot of a desired end state. If, on the other hand, the diversity policy has no independent effect from that of the mitigation policy, then it is an entirely unnecessary policy.

Because of its utilitarian basis, the diversity policy is distinguishable from the end-state approach criticized at the start of this paper. The end-state approach defines justice in terms of an ideal distributional snap-shot. The snap-shot is itself the end toward which the equality monitor strives. Diversity, or in some cases “critical mass”,²¹ is not an end in itself, but a step along the way toward some other goal, depending on the nature of the implementing institution. If diversity achieves the goal for which the policy was implemented, and does so at reasonable expense, then it is welfare enhancing. If one views welfare enhancement as the appropriate goal for the state, then the diversity policy must be applauded when it works to enhance welfare.²² However, if one views justice in distribution as a separate and overriding issue, the diversity policy may be undesirable because it takes us further away from distributional justice even when it is welfare enhancing.

Finally, the argument here implies that equality monitors should be able to defend their actions on moral grounds. The confusion between end-state and process-based notions of equality has generated a bizarre state of affairs, with officials responsible for carrying out mitigation programs presenting poor justifications or, in a surprisingly large number of cases, suggesting that their own actions are morally unjustifiable. To the observer with a generous heart, this can all be attributed to confusion; to the cynic, it is bad faith. There is a moral basis for equality-seeking policies in the racial sphere.

²¹ *Grutter*, at 316.

²² For the argument in favor of welfare enhancement as the appropriate role of government, see Louis Kaplow and Steven Shavell, *Fairness Versus Welfare* (2002).