

BIG DATA AFFIRMATIVE ACTION

Peter N. Salib

ABSTRACT—As a vast and ever-growing body of social-scientific research shows, discrimination remains pervasive in the United States. In education, work, consumer markets, healthcare, criminal justice, and more, Black people fare worse than whites, women worse than men, and so on. Moreover, the evidence now convincingly demonstrates that this inequality is driven by discrimination. Yet solutions are scarce. The best empirical studies find that popular interventions—like diversity seminars and antibias trainings—have little or no effect. And more muscular solutions—like hiring quotas or school busing—are now regularly struck down as illegal. Indeed, in the last thirty years, the Supreme Court has invalidated every such ambitious affirmative action plan that it has reviewed.

This Article proposes a novel solution: Big Data Affirmative Action. Like old-fashioned affirmative action, Big Data Affirmative Action would award benefits to individuals because of their membership in protected groups. Since Black defendants are discriminatorily incarcerated for longer than whites, Big Data Affirmative Action would intervene to reduce their sentences. Since women are paid less than men, it would step in to raise their salaries. But unlike old-fashioned affirmative action, Big Data Affirmative Action would be automated, algorithmic, and precise. Circa 2021, data scientists are already analyzing rich datasets to identify and quantify discriminatory harm. Armed with such quantitative measures, Big Data Affirmative Action algorithms would intervene to automatically adjust flawed human decisions—correcting discriminatory harm but going no further.

Big Data Affirmative Action has two advantages over the alternatives. First, it would actually work. Unlike, say, antibias trainings, Big Data Affirmative Action would operate directly on unfair outcomes, immediately remedying discriminatory harm. Second, Big Data Affirmative Action would be legal, notwithstanding the Supreme Court's recent case law. As argued here, the Court has not, in fact, recently turned against affirmative action. Rather, it has consistently demanded that affirmative action policies both stand on solid empirical ground and be well tailored to remedying only particularized instances of actual discrimination. The policies that the Court recently rejected have failed to do either. Big Data Affirmative Action can easily do both.

AUTHOR—Assistant Professor of Law, University of Houston Law Center; Associated Faculty, University of Houston Hobby School of Public Affairs.

INTRODUCTION 822

I. PERSISTENT DISCRIMINATION AND FAILED INTERVENTIONS 827

 A. *The Science Is Clear: Discrimination Is Everywhere* 828

 B. *Popular Interventions Do Not Work* 832

II. BIG DATA AFFIRMATIVE ACTION 834

 A. *How Big Data Affirmative Action Works* 835

 B. *Big Data Affirmative Action in Action* 839

III. IS BIG DATA AFFIRMATIVE ACTION LEGAL? 849

 A. *The Law of Affirmative Action* 850

 B. *Why Big Data Affirmative Action Passes Muster* 867

 C. *A Goldilocks Problem* 871

IV. NORMATIVE OBJECTIONS 875

 A. *“Reverse Racism” and Signaling Effects* 876

 B. *Individual Injuries, Average Remedies* 878

 C. *Data Skepticism* 881

 D. *Weak Medicine* 883

 E. *Evasion and Adaptive Decision-Making* 884

 F. *Are Humans Obsolete?* 887

 G. *Less Biased Humans or Less Biased Decisions?* 890

CONCLUSION 891

INTRODUCTION

Despite the successes of the Civil Rights Era in rolling back overtly racist laws, and despite Americans’ declining willingness to openly espouse racist views,¹ racial inequality continues to pervade nearly every institution of American life. People of color are arrested and convicted of crimes at higher rates than their white counterparts. Once convicted, they are sentenced to harsher punishment. Educational attainment for nonwhite Americans lags that of white Americans. Statistics are similarly grim for everything from hiring and compensation to healthcare and mortality.² And

¹ There are, of course, notable exceptions to this trend. See, e.g., Richard Fausset, *Rally by White Nationalists Was Over Almost Before It Began*, N.Y. TIMES (Aug. 12, 2018, 3:30 PM), <https://www.nytimes.com/2018/08/12/us/politics/charlottesville-va-protest-unite-the-right.html> [<https://perma.cc/P3SN-XV7R>] (discussing the alt-right movement and white supremacist rallies).

² For a discussion of the vast literature on persistent discrimination across these domains, see *infra* Section I.A.

as artificially intelligent decision-making becomes ascendant, algorithms threaten to replicate and entrench these unequal outcomes permanently.

Solutions are elusive. Empirical studies of popular interventions—like diversity trainings and de-biasing seminars—show that they have little real-world effect. On the other hand, aggressive proposals—like racial hiring quotas—are now regularly struck down by the Supreme Court.

This Article proposes a new paradigm for redressing unfair racial disparities wherever they persist: “Big Data Affirmative Action.” The core idea is simple. Under Big Data Affirmative Action policies, members of racial minority or other disadvantaged³ groups would receive *lighter* sentences, *higher* pay, *better* healthcare, and more. They would receive all of this *because* of their race—offsetting discrimination they would likewise face *because* of race.

But unlike old-fashioned affirmative action plans, Big Data Affirmative Action policies would be carefully calibrated to remedy the exact harm caused by the institutions implementing them. Relying on technologically sophisticated, quantitatively precise empirical foundations, they would distribute benefits large enough to redress measured discrimination, but no larger. In other words, under each policy, the remedy would match the harm. Such careful calibration would allow Big Data Affirmative Action to overcome the legal shortcomings that have recently hamstrung traditional affirmative action interventions like hiring quotas.

This Article proceeds in four Parts. Part I first canvasses the social-scientific evidence on persistent racial inequality. It shows that such inequality continues to plague essentially every important institution of American life. And it argues that recent research clearly demonstrates that discrimination—not other factors—causes this inequality. The Part then examines popular proposed remedies for discriminatory inequality, like de-biasing trainings. It argues that, as far as the best empirical research can show, such strategies are exceedingly unlikely to deliver meaningful results.

Part II then introduces a first-of-its-kind solution: Big Data Affirmative Action. Imagine, for example, a Black job applicant has just been hired at a company with a Big Data Affirmative Action policy for compensation. She and her new employer negotiate a starting salary, which is entered into the company’s HR system. Then an algorithm intervenes, increasing the

³ This Article generally uses racial discrimination as its core example. This is in part for simplicity’s sake and in part because racial discrimination is particularly pervasive. But the ideas and policies discussed here can easily be applied to combat other kinds of discrimination—based on gender, sexual orientation, ethnicity, religion, and more. In addition, where other varieties of discrimination present particular puzzles, these are discussed explicitly. *See, e.g., infra* notes 79–82 and accompanying text (discussing sex discrimination in criminal sentencing).

applicant's pay by precisely the amount the employer likely penalized her because of her race. Or picture instead a Latino defendant who has just been convicted of a crime in a judicial district with Big Data Affirmative Action sentencing. He appears at his sentencing hearing, and the judge hands down a term of imprisonment. Here again, an algorithm steps in. Armed with information about the amount of additional incarceration he incurred because of racial bias in the courthouse, the algorithm reduces his sentence by exactly that amount.

Big Data Affirmative Action has become practicable only recently. As Part II explains, circa 2021, essentially every human decision leaves—or could leave—a data trail. As a result, policymakers can now assemble rich datasets about the multitude of factors influencing decisions ranging from employee hiring to disease diagnosis to criminal sentencing. Such data, paired with sophisticated statistical analysis, can—and indeed, already does—supply new, precise measures of racial discrimination across numerous domains.

In our era of rich, ubiquitous information, the potential for Big Data Affirmative Action is essentially unbounded. Big Data Affirmative Action policies need not be limited to employment or criminal sentencing. They could be implemented wherever people continue to suffer harm from racial bias, so long as sufficient data could be collected to estimate the quantum of that harm. And in our increasingly data-drenched world, that stands to be just about everywhere. In fact, Big Data Affirmative Action could even be applied to *algorithms'* decisions, addressing widespread concerns about algorithmic bias.

Part III turns to Big Data Affirmative Action's legality. Various constitutional and statutory provisions forbid discrimination on the basis of race. And for better or worse, the Supreme Court presently interprets those laws to limit both race-based inflictions of harm and race-based grants of benefit.⁴ This gives rise to the facially puzzling fact that many policies designed to help disadvantaged groups *overcome* discrimination can themselves *constitute* illegal discrimination.

This puzzle is not a new one. Indeed, only half of the Big Data Affirmative Action proposal—the “Big Data” half—is novel. Affirmative action, by contrast, has been the subject of intense legal controversy for over half a century.

Colleges, of course, practice a variety of affirmative action, which has been the subject of particularly heated disagreement in recent years. But campus affirmative action is somewhat idiosyncratic—a minor subgenre of

⁴ See, e.g., *Ricci v. DeStefano*, 557 U.S. 557, 582 (2009).

a formerly grand medium. Today, race is regularly considered in college admissions, but only for the purpose of improving on-campus “diversity”—not redressing discrimination—and never as a decisive factor. This is comparatively weak stuff, not much resembling the Big Data proposal. And at any rate, the Supreme Court has long foreshadowed the demise of diversity-based affirmative action.⁵ It may be poised to deliver on that prediction this Term.⁶

History, however, presents another, more robust species of affirmative action, from which Big Data Affirmative Action more directly descends. In the mid-twentieth century, school busing plans, racial hiring quotas, and set-asides in government spending for minority-owned business all sought to *directly* redress discriminatory harm.⁷ And to achieve that goal, those policies used race as the decisive factor for allocating benefits.

This more muscular variety of affirmative action is widely believed to be either dead or on life support. In the last thirty years, the Supreme Court has struck down every policy of this type that it has reviewed. As a result of those decisions, academics, judges, and lawmakers have repeatedly questioned the continuing legal viability of such affirmative action plans.⁸ Policymakers have thus turned away from affirmative action toward other interventions that are both less legally fraught and, unfortunately, less effective.

Part III goes on to contend that the muscular variety of affirmative action is not dead—only resting. Taking the Supreme Court at its word, the Part argues that this style of affirmative action plan remains legally permissible. But to be permitted, the Court has consistently held, such policies must be empirically sound and carefully designed. As the Court now puts it, affirmative action policies must respond to a “strong basis in evidence” of redressable discrimination.⁹ Part III shows that, when the Court has recently struck down such plans, the reason has been because they failed to clear the required legal hurdle. Big Data Affirmative Action, by contrast, would leap the hurdle easily.

⁵ *Grutter v. Bollinger*, 539 U.S. 306, 343 (2003).

⁶ See *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 142 S. Ct. 895 (2022) (asking the Court to reject diversity-justified affirmative action in college admissions); *infra* Section III.A.1 (explaining the difference between diversity-justified and discrimination-remediating affirmative action).

⁷ See generally *Swann v. Charlotte-Mecklenburg Bd. of Educ.*, 402 U.S. 1 (1971) (busing); *Johnson v. Transp. Agency*, 480 U.S. 616 (1987) (hiring quotas); *Fullilove v. Klutznick*, 448 U.S. 448 (1980) (government spending).

⁸ See *infra* notes 108–110.

⁹ *Wygant v. Jackson Bd. of Educ.*, 476 U.S. 267, 277 (1986).

What exactly is the strong basis in evidence standard, and why would Big Data Affirmative Action satisfy it? The rule has two elements. First, it imposes a burden of proof for showing that actual discrimination is occurring. Crucially, the evinced discrimination must be localized to the institution promulgating the affirmative action policy. An employer, for example, may not attempt to undo discrimination in the education system, nor may educators undo it in the housing system. Each institution may, so to speak, clean only its own house. Second, the rule requires that affirmative action policies be tailored to address only the documented discrimination caused by the promulgating institution—but go no further.

These legal constraints mean each individual affirmative action policy must be careful, imposing modest race-based benefits that will generally fall short of gross racial disparities. For example, a racial hiring gap at a company may be caused in part by the company's own bias against well-qualified Black and Brown candidates. A well-tailored affirmative action policy implemented by the company may redress that quantum of discrimination. But some of the gap may be caused by upstream differences in applicants' attainment of necessary credentials. Even if those differences, too, were caused by discrimination—say, in college admissions—the employer must take them as given. As the strong basis in evidence standard is understood today, an employer's policy that corrects a college's discrimination will be struck down.

Part III then shows why Big Data Affirmative Action policies, if properly designed, would satisfy both elements of the strong basis in evidence standard. Already, careful social scientists are conducting studies that show, with high certainty, how racial bias influences decisions in particular institutions—businesses, industries, schools, courts, and more. These studies show not only that racial bias has *some* effect, but *how much*. They also isolate the sources of discrimination, disentangling the harm attributable to the institution under review from that attributable to other institutions. Big Data Affirmative Action policies would be based on empirical analyses of exactly this kind. They would thus impose race-based adjustments at a particular locus of discrimination, carefully sized to the quantum of discrimination imposed at that locus. As a result—Part III shows—Big Data Affirmative Action's basis in evidence would be even stronger than that of the policies that the Supreme Court has approved.

Part IV turns to normative objections. It begins with the conservative critiques of affirmative action that characterize it as a kind of “reverse racism.” These critiques take at least three forms. One version contends that affirmative action actually harms, not helps, people of color by signaling that they did not earn and do not deserve their achievements. Another version

focuses on the individuals who do *not* benefit from a given affirmative action policy. Here, the argument goes, any race-based benefit to one group inevitably implies an unjust race-based deprivation to another. A third variation takes a more philosophical view, arguing that any policy that treats people differently because of their race is wrong, no matter the consequences.

Part IV argues that these critiques have force—if at all—*only* against old-fashioned, less carefully calibrated forms of affirmative action. Big Data Affirmative Action policies grant precise race-based benefits that remedy what would otherwise be discriminatory harms. They go no further. Thus, beneficiaries do not receive anything undeserved. On the contrary, Big Data Affirmative Action ensures that everyone receives what they *do* deserve—and would have gotten absent discrimination. Likewise, there is no injustice in depriving one group of a benefit secured only via discrimination against another group. Finally, race-based preferences simply are not wrong when they are remedies for race-based discrimination. That is just how remedies work. They flow to the injured.

Part IV addresses other normative questions, too: Are Big Data Affirmative Action remedies sufficiently individualized? Are modern social-scientific measures of discrimination trustworthy enough to base important policies on? Is Big Data Affirmative Action a potent enough intervention to radically reduce racial inequality? Would Big Data Affirmative Action be robust to the dynamic human responses that might follow from its implementation? Should humans continue to have *any* involvement in the making of important, and potentially discriminatory, decisions? Ought we be comfortable with a policy intervention that essentially abandons any hope that humans can change for the better, to be less biased? This Part argues that the answer to all of these questions is “yes.”

I. PERSISTENT DISCRIMINATION AND FAILED INTERVENTIONS

Why do we need Big Data Affirmative Action? Perhaps, one imagines, the worst of America’s racial strife is behind us. And surely whatever discriminatory attitudes remain are soon to be eliminated as antiracism pledges, antibias training, and similar interventions become widespread. Maybe progress is on the march, and the best thing to do is to double down on the strategies that are already spurring it forward.

This Part argues that such optimism, if well-meaning, is deeply misinformed on at least two counts. First, while it is true that the United States’ gravest racial sins—slavery, mass lynchings, Jim Crow—are in the past, inequality persists. In essentially every imaginable domain of life—employment, education, healthcare, criminal justice—Black and brown

Americans suffer worse outcomes than their white peers. And the evidence now shows convincingly that the best explanation for these persistent disparities is widespread discrimination—not other more benign factors.

Second, modern discrimination is stubborn. Antibias trainings and diversity seminars are beginning to pervade public and private institutions. Corporations, universities, and governments profess commitments to antiracism. Yet the best empirical evidence now shows that the leading interventions designed to reduce racism and eliminate inequality in fact have no effect at all. Bolder strategies are needed.

A. *The Science Is Clear: Discrimination Is Everywhere*

American society remains deeply unequal. Worse, the lion's share of disadvantage continues to fall on Americans who, for generations, were disenfranchised, subjugated, dehumanized, and exploited. A vast social science literature investigates this relationship between historical and present-day disadvantage. It catalogues the substantial, measurable disparities that persist today. And crucially, convincing evidence now shows that these disparities are not due to “legitimate” differences between members of different groups. Instead, their cause is discrimination.

Social science skeptics may doubt this claim—or at least its certainty. After all, the world is a messy place, and causal inference is hard. One might readily accept that there are racial, gendered, and other disparities in our society but doubt we have proof that discrimination causes them.

Yet the social-scientific evidence on discrimination is about as rock-solid as it gets. For roughly the past twenty years, economists have been performing randomized experiments that document discrimination across social domains.¹⁰ This style of study design represents the gold standard for causal inference. Bolstering the large and growing body of randomized experimental evidence, innumerable statistical studies using real-world data likewise reveal discrimination. The best of these statistical studies minimize the risk of confounding causes by assembling unusually rich datasets and thus controlling for most or all plausible causal variables.¹¹ Taken together, this substantial corpus of evidence points inexorably in one direction: Race-based inequality, gender-based inequality, and similar social ills are driven by decisions disfavoring members of disadvantaged groups *because of* their membership in those groups.

¹⁰ Marianne Bertrand & Esther Duflo, *Field Experiments on Discrimination*, in 1 HANDBOOK OF ECONOMIC FIELD EXPERIMENTS 309, 310–11, 315–16 (Esther Duflo & Abhijit Banerjee eds., 2017) (collecting studies on discrimination across various social domains).

¹¹ See, e.g., *infra* text accompanying notes 63–71 (discussing Professors M. Marit Rehavi and Sonja Starr's work on criminal sentencing).

A necessarily incomplete survey of the scholarly literature demonstrating pervasive discrimination follows:

- **Discrimination in employment:** Businesses refuse to interview Black workers while readily interviewing otherwise-identical white ones.¹² Employers prefer white applicants over Black ones even in markets for high-skill, college-educated workers.¹³ When employers do hire Black workers, they pay them less than white ones, even when there are no differences in actual productivity.¹⁴ Employers, in fact, discriminate in hiring against *all* nonwhite groups.¹⁵ They also routinely discriminate on the basis of religion, sexuality, age, caste, and gender.¹⁶
- **Discrimination in criminal justice:** Judges and prosecutors impose harsher sentences based on both race and gender. Black and Hispanic defendants are punished more than otherwise-identical white defendants, and men are punished more than women.¹⁷ Prior to sentencing, all-white juries discriminate against Black defendants, convicting them more readily than white defendants.¹⁸ And even

¹² See Marianne Bertrand & Sendhil Mullainathan, *Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination*, 94 AM. ECON. REV. 991, 1011 (2004).

¹³ See John M. Nunley, Adam Pugh, Nicholas Romero & R. Alan Seals, *Racial Discrimination in the Labor Market for Recent College Graduates: Evidence from a Field Experiment*, 15 B.E. J. ECON. ANALYSIS & POL'Y 1093, 1122 (2015).

¹⁴ See Roland G. Fryer Jr., Devah Pager & Jörg L. Spenkuch, *Racial Disparities in Job Finding and Offered Wages*, 56 J.L. & ECON. 633, 635–36 (2013); Valerie Wilson & William M. Rodgers III, *Black-White Wage Gaps Expand with Rising Wage Inequality*, ECON. POL'Y INST. (Sept. 19, 2016), <http://www.epi.org/publication/black-white-wage-gaps-expand-with-rising-wage-inequality/> [<https://perma.cc/JTR8-H8AY>]. In Fryer, Pager & Spenkuch's study, the dataset included workers' wages in their most recent jobs. Fryer et al., *supra*, at 646. This is a nearly direct observation of worker productivity—the foundational “legitimate” trait for which employers are allowed to select when hiring. This sets the study apart from other nonexperimental data-based work. The study found that, conditional on productivity as measured by most recent wages, Black candidates were hired less and paid lower wages than white ones. *Id.* at 635–36.

¹⁵ See Nicolas Jacquemet & Constantine Yannelis, *Indiscriminate Discrimination: A Correspondence Test for Ethnic Homophily in the Chicago Labor Market*, 19 LAB. ECON. 824, 832 (2012).

¹⁶ Bertrand & Duflo, *supra* note 10, at 325–27.

¹⁷ See M. Marit Rehani & Sonja B. Starr, *Racial Disparity in Federal Criminal Sentences*, 122 J. POL. ECON. 1320, 1349 (2014); Crystal S. Yang, *Free at Last? Judicial Discretion and Racial Disparities in Federal Sentencing*, 44 J. LEGAL STUD. 75, 98 (2015); David B. Mustard, *Racial, Ethnic, and Gender Disparities in Sentencing: Evidence from the U.S. Federal Courts*, 44 J.L. & ECON. 285, 298–303 (2001). The extraordinarily rich data and resulting robust controls used in the Rehani & Starr and Yang studies are discussed in Section II.B.1, *infra*.

¹⁸ See Shamena Anwar, Patrick Bayer & Randi Hjalmarsson, *The Impact of Jury Race in Criminal Trials*, 127 Q.J. ECON. 1017, 1032 (2012).

before trial, the police discriminate on the basis of race in searches, citations, and arrests.¹⁹

- ***Discrimination in financial transactions:*** In consumer markets, landlords discriminate on the basis of race, ethnicity, immigration status, and sexual orientation.²⁰ Automobile retailers are less willing to sell to ethnic minorities than to members of the majority group.²¹ Lenders make fewer and worse loans to Black applicants than otherwise-identical white ones.²² Customers disfavor men when asked to tip.²³
- ***Discrimination in healthcare:*** Doctors underestimate the pain experienced by Black patients compared with white patients who report similar pain.²⁴ Emergency medical providers are more likely to transport Black and Hispanic patients to low-quality safety-net emergency rooms than white patients residing in the same zip code.²⁵ Cardiologists are less likely to recommend treatment for Black, Hispanic, and Asian patients, even for serious cardiac disease.²⁶
- ***Discrimination in education:*** In primary and secondary education, teachers and school administrators punish Black children more harshly

¹⁹ See Sarath Sanga, *Reconsidering Racial Bias in Motor Vehicle Searches: Theory and Evidence*, 117 J. POL. ECON. 1155, 1156–59 (2009) (searches); Jeremy West, *Racial Bias in Police Investigations* 19–20 (Oct. 2018) (unpublished manuscript), https://people.ucsc.edu/~jwest1/articles/West_RacialBiasPolice.pdf [<https://perma.cc/F4GZ-QJTA>] (citations); John J. Donohue III & Steven D. Levitt, *The Impact of Race on Policing and Arrests*, 44 J.L. & ECON. 367, 390 (2001) (arrests).

²⁰ Bertrand & Duflo, *supra* note 10, at 328 (collecting studies).

²¹ See Asaf Zussman, *Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars*, 123 ECON. J. F433, F466 (2013).

²² See Devin G. Pope & Justin R. Sydnor, *What's in a Picture? Evidence of Discrimination from Prosper.com*, 46 J. HUM. RES. 53, 65, 89 (2011); see also Ping Cheng, Zhenguo Lin & Yingchun Liu, *Racial Discrepancy in Mortgage Interest Rates*, 51 J. REAL EST. FIN. & ECON. 101, 118 (collecting studies and presenting new results).

²³ See Bharat Chandar, Uri Gneezy, John A. List & Ian Muir, *The Drivers of Social Preferences: Evidence from a Nationwide Tipping Field Experiment* 19 (Nat'l Bureau of Econ. Rsch., Working Paper No. 26380, 2019), https://www.nber.org/system/files/working_papers/w26380/w26380.pdf [<https://perma.cc/8L3P-2S3D>].

²⁴ See Kelly M. Hoffman, Sophie Trawalter, Jordan R. Axt & M. Norman Oliver, *Racial Bias in Pain Assessment and Treatment Recommendations, and False Beliefs About Biological Differences Between Blacks and Whites*, 113 PROCS. NAT'L ACAD. SCIS. 4296, 4296–97 (2016) (collecting studies).

²⁵ See Amresh D. Hanchate, Michael K. Paasche-Orlow, William E. Baker, Meng-Yun Lin, Souvik Banerjee & James Feldman, *Association of Race/Ethnicity with Emergency Department Destination of Emergency Medical Services Transport*, JAMA NETWORK OPEN (Sept. 6, 2019), <https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2749448> [<https://perma.cc/2B2H-BMLG>].

²⁶ See Nancy R. Kressin & Laura A. Petersen, *Racial Differences in the Use of Invasive Cardiovascular Procedures: Review of the Literature and Prescription for Future Research*, 135 ANNALS INTERNAL MED. 352, 363 (2001) (collecting studies).

for a given infraction than their peers.²⁷ And at the postsecondary level, faculty are less responsive to, and less willing to make time for, their female and minority students.²⁸

As a result of this weighty body of evidence, there is no longer much disagreement, even among hard-nosed economists, about whether discrimination is pervasive.²⁹ Instead, the modern academic debate is about *why* people discriminate. Do they engage in “taste-based” discrimination, whereby decision-makers indulge their animus toward certain groups? Or do they instead engage in “statistical discrimination?” There, low-information decision-makers assume each individual member of a disadvantaged group bears some negative trait because the group, on average, is more likely to bear it.³⁰

American antidiscrimination law does not distinguish between these two types of bias. Both are illegal.³¹ That is, it is no more a defense in a discrimination suit to say “I fired him because I assumed that, because he was Black, he would be less productive” than to say “I fired him because I dislike Black people.” The former is forbidden even if, in a given profession, the *average* Black worker could be proven to be less productive than others. There is no legal license to stereotype, even for stereotypes that reflect some kernel of statistical truth.³²

²⁷ See Ying Shi & Maria Zhu, *Equal Time for Equal Crime? Racial Bias in School Discipline* 17 (IZA Inst. of Lab. Econ., Discussion Paper No. 14306, 2021), <http://ftp.iza.org/dp14306.pdf> [<https://perma.cc/27PD-PBH7>].

²⁸ See Katherine L. Milkman, Modupe Akinola & Dolly Chugh, *Temporal Distance and Discrimination: An Audit Study in Academia*, 23 PSYCH. SCI. 710, 711 (2012).

²⁹ See Bertrand & Duflo, *supra* note 10, at 310–12; cf. Jennifer L. Doleac, *A Review of Thomas Sowell's Discrimination and Disparities*, 59 J. ECON. LIT. 574, 576 (2021) (book review) (noting that a rare dissenter from this consensus view did “not engage at all with the large and ever-growing economic literature on whether people are treated differently due to their race”).

³⁰ Bertrand & Duflo, *supra* note 10, at 311.

³¹ *Id.* at 312; *City of L.A. Dep't of Water & Power v. Manhart*, 435 U.S. 702, 708–09 (1978).

³² Outlawing both kinds of discrimination is likely good policy. Historically, some economists suggested that statistical discrimination should be legal since it reflects no animus and could lead to more efficient allocations of labor. Stewart J. Schwab, *Is Statistical Discrimination Efficient?*, 76 AM. ECON. REV. 228, 228 (1986). However, more recent scholarship suggests that this is wrong. It shows how statistical discrimination in fact benefits individual institutions at the expense of society as a whole. See Doleac, *supra* note 29, at 577–79 (collecting studies). In short, statistical discrimination rationally disincentivizes those at whom it is aimed from investing in valuable skills. This increases statistical discrimination in turn, setting off a downward spiral. Thus, in addition to being morally objectionable, statistical discrimination is likely inefficient at the societal level. Our legal rules forbidding statistical discrimination therefore solve a collective action problem, making everyone, including employers, better off.

In sum, the evidence is clear. Discrimination persists. It produces substantial inequality, pervading all aspects of everyday life. Solutions are sorely needed.

B. Popular Interventions Do Not Work

Assuming modern discrimination is driven by widespread prejudicial attitudes—whether conscious or unconscious—one appealing solution would be to eliminate those attitudes. If there were a reliable treatment to cure discriminatory thinking, or even just reduce it substantially, one might hope that racial equality would follow. And of course, ridding people of their invidious attitudes surely has intrinsic value as well.

Proposals for “de-biasing” treatments therefore abound, including from respected empirical psychologists. These proposals are creative and diverse. Researchers advocate: “cognitive conditioning” (training the mind to regard minority groups positively); “contact” interventions (promoting interaction with members of minority groups); individual introspection, including on moral values like equality; and more.³³ Today, such treatments are regularly packaged into innumerable variegated “diversity trainings” or “antibias seminars” and deployed at corporate headquarters, university campuses, public schools, and beyond.

De-biasing treatments are not the only possible mainstream alternative to Big Data Affirmative Action. Another intervention for curing discrimination might be to blind more decision-makers to their subjects’ race. This could work in some contexts, but the applications are limited. It is hard to imagine, for example, a judge sentencing a criminal defendant without ever meeting him, even by video. The same goes for hiring. Few companies are willing to make long-term hiring decisions without requiring even a single meeting from which race could be inferred. De-biasing trainings and Big Data Affirmative Action are thus natural comparators in that they are ambitious in their scope. Both solutions aim for broad applicability.

The problem with de-biasing trainings is that they do not actually work. A meta-study by Elizabeth Levy Paluck, Roni Porat, Chelsey Clark, and Donald Green, published last year in the *Annual Review of Psychology*, evaluates the state of the evidence. Paluck and her colleagues performed a “random effects meta-analysis of all prejudice-reduction interventions from

³³ Elizabeth Levy Paluck, Roni Porat, Chelsey S. Clark & Donald P. Green, *Prejudice Reduction: Progress and Challenges*, 72 ANN. REV. PSYCH. 533, 543–45, 547 (2021).

2007 to 2019.”³⁴ They evaluated some “418 studies, which report 1,292 distinct point estimates.”³⁵

The evaluated studies themselves report only modest effects. On average, the studies claim that their interventions are only good enough to move someone with a “mildly negative feeling” about Black people to “nearly reaching a neutral feeling.”³⁶ Another “way to look at this effect size is that following an intervention, there would be an 85% overlap between people who went through the intervention and people who were in the control group.”³⁷ Moreover, these results are mostly from the lab. When interventions are deployed in real-world settings like diversity trainings, the studies’ reported effects shrink even further.³⁸

Also worth noting is the fact that almost none of the literature even attempts to measure the effect of antibias interventions on discriminatory *behavior*. Most studies instead measure the interventions’ effects on conscious or unconscious racial *attitudes*.³⁹ But biased attitudes do not equal biased behavior, and the link between unconscious bias and behavior is especially controversial.⁴⁰ Furthermore, even assuming attitudes did strongly affect behavior, the literature reports that de-biasing interventions’ effects on attitudes are fleeting. They fade to nothing after just a few hours or, at most, a few days.⁴¹

Worse, there is substantial evidence that these modest results reported in the primary literature are vastly overblown. De-biasing research appears to be yet another victim of the “replication crisis” sweeping across empirical psychology as a whole.⁴² One source of this widespread nonreplicability is “publication bias”—whereby splashy, but statistically random, results get published, while null results get trashed. As Paluck and her coauthors write, “A telltale sign of publication bias is a strong positive relationship between reported effects and their standard errors Our collection of studies

³⁴ *Id.* at 539.

³⁵ *Id.*

³⁶ *Id.* at 540.

³⁷ *Id.* at 553.

³⁸ *Id.* at 542–43.

³⁹ *Id.* at 548–49.

⁴⁰ *Id.* at 548–49, 553.

⁴¹ *Id.* at 550, 545–46; Calvin K. Lai et al., *Reducing Implicit Racial Preferences: II. Intervention Effectiveness Across Time*, 145 J. EXPERIMENTAL PSYCH. 1001, 1012 (2016).

⁴² Bradford J. Wiggins & Cody D. Christopherson, *The Replication Crisis in Psychology: An Overview for Theoretical and Philosophical Psychology*, 39 J. THEORETICAL & PHIL. PSYCH. 202, 202–03 (2019); Paluck et al., *supra* note 33, at 538.

displays a powerful relationship of this kind.”⁴³ As a result, the effect-size results “are not robust to the most basic assessments of study quality.”⁴⁴

“In the absence of publication bias, we should obtain similar average effect estimates from small and from large studies.”⁴⁵ Yet in the de-biasing literature, small studies—with less statistical power—consistently reported substantially larger effects than large ones. On average, the smallest studies reported more than *twice* the effect of the largest. When only large, high-powered studies are considered, the effect size is minute.⁴⁶ Indeed, the publication bias in this literature appears so strong that

a study large enough to generate a standard error of approximately zero would, on average, produce no change in prejudice at all. In other words, if the current collection of studies had been conducted on a much larger scale, our analysis would have shown *no reduction in prejudice*.⁴⁷

Unfortunately, we do not yet know how to reliably make humans meaningfully less racist, sexist, or otherwise prejudiced. Even the rosier evaluations of our best interventions show modest results—far smaller than would be needed to meaningfully redress discrimination in any domain. And those rosy estimates are quite likely to be wildly overconfident. This does not mean that researchers will never figure out how to fix discrimination from the inside out. But as of today, we cannot do it, and interventions designed around the idea of de-biasing human attitudes or behaviors will not work. New proposals are needed now. The millions of Americans today who are systematically shunted into worse jobs, increased criminalization, less healthcare, and other unfair outcomes cannot afford to wait.

II. BIG DATA AFFIRMATIVE ACTION

Affirmative action, unlike other policy interventions, *does* work. “A long literature documents” the success of old-fashioned affirmative action policies in reducing unfair outcomes for disadvantaged groups.⁴⁸ Such policies have, for example, opened the door to Black and female workers in industries where they were previously unwelcome.⁴⁹ They have diversified

⁴³ Paluck et al., *supra* note 33, at 538.

⁴⁴ *Id.* at 541.

⁴⁵ *Id.*

⁴⁶ *Id.* at 553–54.

⁴⁷ *Id.* at 538 (emphasis added).

⁴⁸ Doleac, *supra* note 29, at 584–85.

⁴⁹ *Id.* at 584; Fidan Ana Kurtulus, *The Impact of Affirmative Action on the Employment of Minorities and Women: A Longitudinal Analysis Using Three Decades of EEO-1 Filings*, 35 J. POL’Y ANALYSIS & MGMT. 34, 39–40 (2016).

police forces.⁵⁰ In fact, affirmative action is powerful enough that its discrimination-reducing effects persist after the intervention has lapsed.⁵¹

Big Data Affirmative Action is the next evolutionary phase of this tried-and-true approach. It is affirmative action in the traditional sense: it disburses benefits to members of disadvantaged groups (racial minorities, religious minorities, sexual minorities, women) *because* of their membership in those groups. But unlike old-fashioned affirmative action, Big Data Affirmative Action distributes benefits automatically, algorithmically, and precisely.

This Part describes the technology underlying Big Data Affirmative Action and explains how it could be used as a general matter to correct wrongful discrimination. It then proposes specific real-world examples.

A. How Big Data Affirmative Action Works

Big Data Affirmative Action operates via sophisticated statistical analysis of large datasets obtained and maintained by the potential discriminator: employers, admissions offices, hospitals, courts, and so on. The analysis produces measurements that determine both *who* gets benefits and *how much* is due.

Large employers today, for example, keep records of hundreds of thousands of job applications, hiring decisions, salary offers, and promotions. These datasets are not simply large; they are *rich*. They contain information about workers' educational history, employment history, geography, and more.⁵² They also can, and often do, contain information about workers' membership in legally protected classes—like race, gender, or sexual orientation.⁵³ Employers further enrich these datasets over the course of the hiring process, recording job-relevant performance metrics like test scores and peer reviews.⁵⁴ The resulting datasets contain information

⁵⁰ Doleac, *supra* note 29, at 584 n.7; Amalia R. Miller & Carmit Segal, *Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement*, 94 REV. ECON. & STATS. 1107, 1124 (2012).

⁵¹ Doleac, *supra* note 29, at 584; Conrad Miller, *The Persistent Effect of Temporary Affirmative Action*, 9 AM. ECON. J. 152, 153 (2017).

⁵² See, e.g., Elena Yakimova, *6 Steps to Make Your Recruiter Database Better*, SPICEWORKS (Dec. 16, 2021), <https://www.spiceworks.com/hr/recruitment-onboarding/articles/6-steps-to-make-your-recruiter-database-better/> [<https://perma.cc/TKC2-EFUV>] (describing how human-resources departments can use a database-management system to manage recruiter data, including candidates' current and past employment, marital status, and more).

⁵³ See *National Origin Discrimination – FAQs*, U.S. EQUAL EMP. OPPORTUNITY COMM'N, <https://www.eeoc.gov/youth/national-origin-discrimination-faqs> [<https://perma.cc/J4B5-A7GF>] (noting that employers collect such information, among other reasons, to satisfy reporting requirements).

⁵⁴ See Dana Pessach, Gonen Singer, Dan Avrahami, Hila Chalutz Ben-Gal, Erez Shmueli & Irad Ben-Gal, *Employees Recruitment: A Prescriptive Analytics Approach via Machine Learning and Mathematical Programming*, 134 DECISION SUPPORT SYS. 1, 6 (2020).

about essentially every factor that, rightly or wrongly, might have influenced each ultimate employment decision.

Armed with such rich data, careful practitioners of modern statistical analysis can tease out which factors actually did influence the final decisions. They can do so using straightforward and well-understood statistical models like least squares regressions. Regression models work by isolating the distinct relationship that each of many different inputs has on a single output of interest.

To create a Big Data Affirmative Action plan, such a model would first be used to determine, for example, the effect of job applicants' race on their salary. Using a regression model or a similar technique, data scientists would "hold constant" the other explanatory variables in the data, thus isolating racial penalties. Put another way, they would determine whether, once "legitimate" factors influencing pay (education, experience, test scores, etc.) were held equal, race-correlated salary disparities remained. If so, the best inference would be that discrimination was the cause.

This kind of statistical analysis is by no means theoretical. Such data work has long been the bread and butter of empirical social science; innumerable examples exist.⁵⁵ Indeed, regression analysis and similar approaches are so well understood as to be somewhat old-fashioned. Recent years have seen an explosion in the use of new machine learning models—like neural networks or random forests—to make decisions based on large datasets.⁵⁶ However, such newer models would generally not be ideal for use in Big Data Affirmative Action. Unlike regression models, these new tools are not designed to cleanly isolate the effect of various inputs on a given output.⁵⁷ Thus, while useful for many tasks, they are not well suited for analyzing a large dataset to isolate the effect of race—or another protected characteristic—on a particular outcome.

Astute readers will note here that randomized controlled trials, not statistical analyses, are the gold-standard technique for identifying cause and effect. True. But as just discussed, many such trials have *already* established that discrimination is pervasive, driving disparate outcomes across social domains. This should update our priors significantly in favor of data-analytic analyses that show the same. It is one thing to be skeptical of statistics when a general effect is not well established. It is another entirely to be skeptical when the statistics simply confirm *yet another* instance of a well-established phenomenon. Not even careful academic econometricians would argue that,

⁵⁵ *Supra* Section I.A (describing studies on race-based discrimination across social domains).

⁵⁶ TREVOR HASTIE, ROBERT TIBSHIRANI & JEROME FRIEDMAN, *THE ELEMENTS OF STATISTICAL LEARNING: DATA MINING, INFERENCE, AND PREDICTION* 9, 337, 352, 412 (2d ed. 2009).

⁵⁷ See Peter N. Salib, *Artificially Intelligent Class Actions*, 100 *TEX. L. REV.* 519, 543 (2022).

once a general problem is rigorously established, every individual policy response to that problem requires original research publishable in *Econometrica*. Thus, as policy interventions go, Big Data Affirmative Action's grounding in statistical analysis is about as scientific as it gets.⁵⁸

Importantly, the statistical models undergirding Big Data Affirmative Action would reveal not only the existence but also the quantum of discriminatory harm. Once an employer knows, for example, how much less money its Black employees are paid *on account of their being Black*, the employer can do something about it. Under a Big Data Affirmative Action plan, an employer would adjust salary offers for Black employees upward by just that amount. Such adjustments could be made algorithmically and automatically. Neither employees nor employers would have to make any conscious changes. Everyone would proceed normally through the job application and hiring process, including salary negotiations. Upon entry of the new hire's salary into an HR system, an algorithm could simply increase that salary by the amount the statistical model prescribed. Similar adjustments could be made algorithmically for existing Black employees, too.

Moreover, Big Data Affirmative Action policies could correct for multiple kinds of discrimination at once. Suppose a job applicant is both Black and a woman. Both groups face pay discrimination on account of their identities. But it does not necessarily follow that the discrimination is additive and linear. That is, the penalty for being a Black woman may not simply be the penalty for being a woman plus the penalty for being Black. A sophisticated Big Data Affirmative Action plan could include an analysis of—and thus a correction for—discrimination against Black women, specifically. In this way, Big Data Affirmative Action can be sensitive to what is sometimes called the “intersectionality” of discrimination.⁵⁹

How large would the benefits disbursed under Big Data Affirmative Action be? This would vary by context, since the amount of discrimination would likewise vary with the institution implementing the policy. It is important to note, however, that any single Big Data Affirmative Action policy's benefits would almost inevitably be *smaller* than gross observed

⁵⁸ Big Data Affirmative Action need not always rely on data analytics. In some cases, discrimination might be measured using true randomized controlled trials—as with the economic studies that randomly vary names on resumes. See Bertrand & Duflo, *supra* note 10, at 319. However, such studies are of limited use when the decision under examination involves direct human-to-human contact. *Id.* at 318. There is, as of yet, no way to randomly assign a race, gender, or other protected identity to a living breathing person. Thus, many—likely most—Big Data Affirmative Action policies would be grounded in statistical analysis of real-world data.

⁵⁹ Kimberle Crenshaw, *Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics*, 1989 U. CHI. LEGAL F. 139, 140.

racial or other disparities. Consider again an employer that wishes to adjust salaries using Big Data Affirmative Action. Suppose that it observes that its Hispanic employees are paid, on average, 30% less than their white colleagues. At first blush, then, it seems that the company should adjust every Hispanic employee's pay upward by thirty cents on the dollar.

But perhaps the mix of jobs within these two groups is different. Perhaps white employees are comparatively likely to work as engineers, whose skills are particularly scarce and expensive. Moreover, even within job functions, there might be legitimate differences. Perhaps among engineers, white workers are more likely to be proficient in Python, a highly in-demand programming language.⁶⁰ This, too, might be a justifiable reason for additional pay, assuming that Python skills remain scarce relative to demand. A Big Data Affirmative Action policy would control for such legitimate differences in job function. As a result, the policy's salary adjustments would account only for the differences between, for example, Hispanic and white Python programmers. Once non-race-based, legitimately-work-related factors were accounted for, the observed Hispanic-white disparity, and thus the Big Data Affirmative Action adjustment, might only be, say, 10%.

Note that well-designed Big Data Affirmative Action policies could and should be updated regularly. Either the implementation of such a policy or secular trends could cause measured discrimination to change over time. Thus, if subsequent data showed a shift in discrimination over time, an updated policy might increase Hispanic engineers' salaries by, say, 8 or 12%, rather than 10%.

It is tempting to insist on a salary adjustment of the full 30% disparity for the company's Hispanic workers. This is especially true because the aforementioned "legitimate" differences in job functions or qualifications might themselves be caused by discrimination. Perhaps more white than Hispanic employees are software engineers because of discriminatory practices in funding and allocating elementary math instruction. Perhaps differences in technical qualifications among engineers were caused by similar race-based exclusion in college or graduate education. Yet a well-designed Big Data Affirmative Action policy would ignore all of this potential upstream discrimination.

This conservative approach is pragmatic. As will be discussed below, the legal rules governing affirmative action policies are strict. And an affirmative action policy that is struck down as illegal helps no one. Thus, a

⁶⁰ See *11 Most In-Demand Programming Languages in 2022*, BERKELEY BOOT CAMPS, <https://bootcamp.berkeley.edu/blog/most-in-demand-programming-languages/> [<https://perma.cc/D5LY-YP2N>].

cardinal objective in designing Big Data Affirmative Action policies is to comply with the law. And as it stands, the law currently limits a given institution's affirmative action policies to correcting only the institution's *own* discrimination.⁶¹ A company may correct its own failure to pay Hispanic engineers the same as similarly qualified white engineers. But it has no legal authority to correct for discrimination in primary education, college, or elsewhere.

Nonetheless, Big Data Affirmative Action could be a powerful tool for redressing discrimination. Certainly, each individual policy would correct only a modest quantum of discrimination imposed by a single institution. But as discussed below, if such policies were implemented by many institutions, they would aggregate to produce large effects.

There is ample reason to think that Big Data Affirmative Action could be implemented broadly by various institutions whose decisions drive persistent racial inequality. Lots of them—employers, hospitals, courts, lenders—already have the necessary data. From there, the task is simple: take the existing data, hire a competent data scientist, and turn her loose. This, too, is already being done in some corners of the private sector. Uber, for example, recently hired the eminent economist John List, who used its rich datasets covering millions of drivers and tens of millions of rides to determine whether Uber drivers faced a discriminatory gender-based pay gap.⁶² Nothing stops more institutions from doing the same.

The potential applications of Big Data Affirmative Action are far too numerous to explore exhaustively here. However, by way of survey, the next Section sketches three specific examples of Big Data Affirmative Action policies: one for criminal sentencing, one to correct employment discrimination, and one to combat the much-ballyhooed threat of algorithmic discrimination. These examples highlight both the diversity of potential applications for Big Data Affirmative Action, as well as possible challenges. The latter, it will be argued, can generally be overcome.

B. *Big Data Affirmative Action in Action*

This Section describes three example Big Data Affirmative Action policies: for discrimination in criminal sentencing, discrimination in pay,

⁶¹ See *infra* note 166 and accompanying text.

⁶² See Alison Griswold, *Uber's Secret Weapon Is Its Team of Economists*, QUARTZ (July 20, 2022), <https://qz.com/1367800/ubernomics-is-ubers-semi-secret-internal-economics-department/> [<https://perma.cc/4SKZ-HCY9>]; Cody Cook, Rebecca Diamond, Jonathan V. Hall, John A. List & Paul Oyer, *The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers* 1 (Nat'l Bureau of Econ. Rsch., Working Paper No. 24732, 2018), https://www.nber.org/system/files/working_papers/w24732/w24732.pdf [<https://perma.cc/LE2T-JK67>].

and algorithmic discrimination. These examples demonstrate Big Data Affirmative Action's broad potential to roll back unfair discrimination across a wide range of distinct and important contexts. The goal here is to give a sense of feasibility. Each context presents its own unique hurdles. But despite their differences, each set of challenges to implementing Big Data Affirmative Action can be overcome.

1. De-Biasing Criminal Sentences

Criminal sentences are wildly unequal. In raw terms, Black defendants are sentenced to terms of imprisonment between 50% and 100% longer than whites.⁶³ Recently, academics have sought to quantify the proportion of this disparity attributable to discrimination in the courthouse. In doing so, they have produced exactly the kind of statistical evidence on which Big Data Affirmative Action policies could be based.

In their recent paper, Professors M. Marit Rehavi and Sonja B. Starr estimate the magnitude of discrimination by members of the federal criminal justice system. Rehavi and Starr leverage a large "linked multiagency data set that follows federal cases from arrest through to sentencing."⁶⁴ This includes "individual records from the US Marshals Service (which collects arrest data), federal prosecutors, federal courts, and the US Sentencing Commission," constituting "a complete picture of each individual's path through the federal justice system."⁶⁵ The resulting dataset catalogues the information of 36,659 men arrested between 2006 and 2008.⁶⁶

Such rich data allow Rehavi and Starr to control for essentially every credible "legitimate" factor that might explain the racial disparity. They use U.S. Marshals Service records to control for the arrest offense, as recorded by the arresting officer.⁶⁷ This helps to ensure that estimates of racial discrimination are based on actual criminal conduct rather than bargained (or biased) charging decisions. They control for the presence of multiple defendants—which indicates a conspiracy and more serious criminal conduct. They control for criminal history, since serial offenders may justifiably be punished more harshly. They control for possible differences in effectiveness between publicly appointed and privately retained counsel. They even control for demographic factors like education, age, geography, income, and employment.⁶⁸ The wisdom of basing criminal sentences on

⁶³ Rehavi & Starr, *supra* note 17, at 1337–38.

⁶⁴ *Id.* at 1321.

⁶⁵ *Id.*

⁶⁶ *Id.* at 1328.

⁶⁷ *Id.* at 1331, 1337. Rehavi and Starr use multiple approaches to capturing the severity of actual conduct, including ones designed to account for possible gaps in the data. *Id.* at 1338–43.

⁶⁸ *Id.* at 1331, 1337, 1342.

such demographic factors is questionable, to say the least.⁶⁹ But variation along these lines is at least formally race neutral.

Even controlling for all of these factors, Rehavi and Starr find that Black defendants are sentenced to terms of incarceration 9% longer than otherwise-identical white defendants.⁷⁰ Using an alternate statistical approach but the same controls, they reach a similar result: Black defendants are sentenced to an average of seven months more incarceration than whites—a gap explained only by race.⁷¹

Other similar studies have reached much the same conclusion. Professor Crystal S. Yang has estimated that federal judges use their discretion to impose three additional months of imprisonment on Black defendants.⁷² This is consistent with Rehavi and Starr's finding that roughly half of their seven-month disparity may be attributable to prosecutorial discrimination, and the other half judicial.⁷³ Writing a decade before Yang, Starr, and Rehavi, Professor David B. Mustard found that Black defendants were sentenced to 4.81 extra months of incarceration compared with whites.⁷⁴ This estimate is congruent with Yang's evidence that the Supreme Court's decision in *United States v. Booker*, decided after Mustard's study but before hers, increased the Black–white gap by about two months.⁷⁵

Studies like these—both their findings and their mere existence—show three related things. First, there is substantial and empirically verifiable discrimination happening in federal courthouses, driving racial disparities in sentencing. Second, discovering and quantifying such discrimination is possible. The data can be assembled, even if doing so requires linking different agencies' databases. True, the above studies are of federal criminal defendants, and states have long lagged behind the feds in their data collection practices.⁷⁶ But this has begun to change in recent years, with nineteen states having recently mandated more-comprehensive data

⁶⁹ Reliance on such factors was mostly disallowed before *United States v. Booker*, 543 U.S. 220 (2005), and remains disfavored. See Yang, *supra* note 17, at 98.

⁷⁰ Rehavi & Starr, *supra* note 17, at 1338.

⁷¹ *Id.*

⁷² Yang, *supra* note 17, at 108.

⁷³ Rehavi & Starr, *supra* note 17, at 1343.

⁷⁴ Mustard, *supra* note 17, at 298.

⁷⁵ 543 U.S. 220 (2005). *Booker* made the Federal Sentencing Guidelines optional, increasing the opportunity for discrimination in sentencing. Yang, *supra* note 17, at 77.

⁷⁶ See Amy Bach, *Missing: Criminal Justice Data*, N.Y. TIMES (Mar. 22, 2018), <https://www.nytimes.com/2018/03/21/opinion/missing-criminal-justice-data.html> [<https://perma.cc/R2XA-PNPP>].

collection.⁷⁷ This suggests that implementing Big Data Affirmative Action in state criminal courts is largely a question of will. States that wish to do it can mandate both the affirmative action and the data collection needed to implement it.

Third and finally, these studies show that Big Data Affirmative Action is a viable policy tool for correcting discrimination in criminal sentencing. Statistical analyses like Rehavi and Starr's can supply the information needed to correct biased courthouse decisions. These studies are precise. As already noted, Rehavi and Starr found that the penalty for being Black in the federal criminal justice system is 9% more incarceration. The Big Data Affirmative Action remedy, then, is to *reduce* Black defendants' sentences by 9%. The literature also estimates discrimination—and thus quantifies the necessary Big Data Affirmative Action correction—for other groups. Yang's results show that Hispanic defendants receive sentences, on average, 1.9 months longer than similar white defendants.⁷⁸ Under a Big Data Affirmative Action sentencing policy, their sentences would then be reduced by 1.9 months on average.

According to Mustard's analysis, women's sentences are over five months shorter than comparable men's sentences.⁷⁹ This result raises an interesting puzzle for Big Data Affirmative Action: In the presence of discrimination, which outcome should be viewed as the "correct" baseline? In the case of, for example, Black–white disparities, it is natural to assume that white defendants are being treated fairly and Black defendants are being unjustifiably mistreated. That inference is easy, given the hundreds of years of outrageous, yet widely accepted, mistreatment Black Americans have suffered. But in the case of male–female disparities in incarceration, it is tempting to assume that female defendants are being treated irrationally *well*. It is yet another familiar, pernicious stereotype that women are passive, weak, and impulsive. Such tropes could add up to a misperception of women as less dangerous and thus less worthy of punishment than an objective assessment would reveal.

What to do? Should Big Data Affirmative Action sentencing policies increase women's terms of incarceration to match men's? Should it reduce men's incarceration to match women's? Should it do a bit of both, so that the genders meet in the middle? Substantive antidiscrimination law generally

⁷⁷ See *Examples of Statewide Criminal Justice Data Repositories*, OHIO CRIM. SENT'G COMM'N (Dec. 2020), <https://www.supremecourt.ohio.gov/Boards/Sentencing/committees/uniformSentEntry/RepositoryExamplesFAQ.pdf> [<https://perma.cc/NLK8-AYND>] (collecting statutes).

⁷⁸ Yang, *supra* note 17, at 92.

⁷⁹ Mustard, *supra* note 17, at 298.

permits any of these as a cure for discrimination.⁸⁰ Thus, in most Big Data Affirmative Action contexts, the answer to this kind of question will depend on one’s normative—or perhaps empirical—views. Here, the “correct” level of criminal sanction will depend on one’s theory of punishment, along with the evidence about the best way to achieve a given theory’s goals. And once that “correct” sentence is determined, a Big Data Affirmative Action algorithm can intervene to adjust sentences accordingly to match the desired output.

In the sentencing context, *Booker* complicates things—but only slightly. In that case, the Supreme Court held that, under the Sixth Amendment, a jury must find any facts that are used to increase incarceration.⁸¹ Thus, a Big Data Affirmative Action policy that increased women’s sentences would have to include jury findings about the operative information. This would not usually be difficult, since proving a defendant’s gender beyond a reasonable doubt would almost always be trivial.⁸² However, a Big Data Affirmative Action policy that only *decreased* sentences would face no *Booker* problem at all.

2. Correcting Pay Disparities

Big Data Affirmative Action for pay disparities presents comparatively few complications. As with criminal sentencing, voluminous evidence from social science reveals that women and minorities are paid less than relevantly identical white men. Much of this research proceeds, like the criminal sentencing studies, via statistical analysis of rich datasets. Professors Kevin Lang and Michael Manove, for example, find a substantial market-wide Black–white wage gap, even after controlling for cognitive aptitude, level of education, quality of education, age, parents’ education, number of siblings, immigration status, parents’ immigration status, and geography.⁸³ Many additional data-analytic studies likewise reveal invidious disparities

⁸⁰ *Sessions v. Morales-Santana*, 137 S. Ct. 1678, 1698 (2017) (describing the “remedial alternatives”).

⁸¹ *United States v. Booker*, 543 U.S. 220, 244 (2005). Technically, *Booker* holds that juries must find any facts that would automatically increase the maximum sentence. *Id.* at 227–28. But the logic applies straightforwardly to a Big Data Affirmative Action policy that automatically increases the defendant’s *actual* sentence.

⁸² The particular facts affecting a sentence would depend on the statistical evidence of discrimination. If gender interacted with age in some way—say, if young women got even more favorable treatment than older ones—both facts would have to be proven. In any case, this kind of demographic information will be controversial only rarely. Indeed, in the mine-run of cases, a defendant’s sustained denial of a long-held identity or easily proved fact would likely be sanctionable. 18 U.S.C. § 3162(b)(2).

⁸³ Kevin Lang & Michael Manove, *Education and Labor Market Discrimination*, 101 AM. ECON. REV. 1467, 1482, 1491 (2011).

unexplained by “legitimate” wage determinants.⁸⁴ Other studies are experimental. Professors Marianne Bertrand and Sendhil Mullainathan famously show that employers are 50% more likely to respond to resumes bearing white-sounding names than to otherwise-identical resumes carrying Black-sounding names.⁸⁵ Still other studies take even more creative approaches. Professors Roland Fryer, Devah Pager, and Jörg Spenkuch compare employers’ wage offers to Black and white job seekers to those workers’ previous salaries, using the latter figure as a metric of actual productivity. They, too, find substantial racial discrimination.⁸⁶

Worth noticing here is not just the fact that economists can—and do—generate highly credible estimates of pay discrimination in various segments of the labor force. It is that they do so despite their task being much harder than it would be for employers to measure their *own* discrimination. Economists are generally on the outside looking in. They therefore have to make creative guesses about what information is available to employers offering wages and, moreover, about what information employers actually considered.⁸⁷ Employers, by contrast, know all of that. They know what information is collected in the application process. They know what questions are asked in the interviews. They know the metrics on which applicants and current employees are evaluated.

Insofar as employers do not know any of this, they can find out, including by mandate. That is, employers can simply decide as a matter of

⁸⁴ See generally, e.g., Mary Corcoran & Greg J. Duncan, *Work History, Labor Force Attachment, and Earnings Differences Between the Races and Sexes*, 14 J. HUM. RES. 3 (1979) (finding that differences in education, training, and job performance explained only a small proportion of the wage gaps between white men and other groups); Francine D. Blau & Andrea H. Beller, *Black-White Earnings over the 1970s and 1980s: Gender Differences in Trends*, 74 REV. ECON. & STAT. 276 (1992) (finding unexplained differences and tracking their variation over time); Ronald L. Oaxaca & Michael R. Ransom, *On Discrimination and the Decomposition of Wage Differentials*, 61 J. ECONOMETRICS 5, 12–18 (1994) (employing a variety of methods to show the Black–white wage gap is explained at least in part by labor-market discrimination); William A. Darity Jr. & Patrick L. Mason, *Evidence on Discrimination in Employment: Codes of Color, Codes of Gender*, 12 J. ECON. PERSPS. 63 (1998) (collecting studies on employment discrimination); Cordelia W. Reimers, *Labor Market Discrimination Against Hispanic and Black Men*, 65 REV. ECON. & STAT. 570 (1983) (finding that differences in measured characteristics, not overt discrimination, overwhelmingly explained the white–Hispanic wage gap).

⁸⁵ Bertrand & Mullainathan, *supra* note 12, at 997–98; see also, e.g., Michael Firth, *Racial Discrimination in the British Labor Market*, 34 INDUS. & LAB. RELS. REV. 265, 272 (1981) (finding considerable discrimination in British employers’ assessment of job applicants of various nationalities); A. Esmail & S. Everington, *Racial Discrimination Against Doctors from Ethnic Minorities*, 306 BRITISH MED. J. 691, 692 (1993) (finding applicants with English names were twice as likely to be selected for hospital posts than candidates with Asian names).

⁸⁶ See Fryer et al., *supra* note 14, at 634–36.

⁸⁷ See Bertrand & Duflo, *supra* note 10, at 332 (“[W]ith more continuous outcome variables—ones that typically are not available to the researcher, such as the ranking of the job candidates by the employer—[it] would . . . be possible to resolve this tension.”).

policy what does and does not matter for hiring or setting pay. In this way, they can—and already do—limit what hiring managers may permissibly consider. Then employers can collect *that* information and hand the data over to statistical empiricists.⁸⁸ In the resulting study, the institutionally approved job-relevant criteria would function as controls. And any unexplained race-correlated disparities not eliminated by those controls would count as discrimination by the firm. The measured disparities, in turn, would be used to design a Big Data Affirmative Action algorithm that eliminated them *ex post*.

One could argue that this approach fails to produce true measures of employer discrimination. Perhaps some hiring managers would defy company policy and consider additional race-correlated factors that, in their views, mattered to job performance. Those would arguably be “legitimate” factors explaining pay disparities but not controlled for in the statistical analysis. On the other hand, a hiring manager’s *ad hoc* reliance on unapproved, race-correlated attributes might just as likely mask the manager’s own racial bias as it may accurately predict performance. By comparison, firm-level evaluations of performance metrics rest on substantially more data and experience, both across the firm and over time, than those of a single manager.

Perhaps for these reasons, when a firm decides to designate certain factors as irrelevant to job performance and pay, antidiscrimination law treats that policy as decisive.⁸⁹ A business cannot rebut, for example, a compensation policy’s racially disparate impact by arguing that the disparity is explained by factors it has affirmatively *disavowed* as relevant.⁹⁰ It is therefore both normatively and legally appealing to treat race-correlated salary disparities unrelated to a company’s carefully considered list of job qualifications as remediable discrimination.

3. *Undoing Algorithmic Discrimination*

Algorithmic bias is a topic *du jour* in the antidiscrimination literature. Algorithms—both the simple and the bafflingly complex—now guide decisions across an incredible range of domains.⁹¹ They make hiring recommendations, determine who gets bail and who is jailed,

⁸⁸ See *supra* note 62 and accompanying text.

⁸⁹ *Ricci v. DeStefano*, 557 U.S. 557, 564–67, 587–88 (2009) (treating an employer’s official policy regarding job-related skills as decisive in a disparate impact case).

⁹⁰ See *id.* at 587–89.

⁹¹ See Salib, *supra* note 57, at 567.

authorize or deny home loans, and more.⁹² Many scholars argue that the use of such algorithmic decision-making causes or entrenches discriminatory outcomes.⁹³

Some algorithmic discrimination can be mitigated via careful algorithmic design.⁹⁴ But in the end, algorithms that mimic human decisions risk mimicking human bias too. Consider, for example, an algorithm trained to automate a firm's hiring recommendations based on data from its past hires. Here, certain input data may be biased. Nonwhite applicants may be more likely to have attended for-profit colleges not ranked by *U.S. News*.⁹⁵ If the hiring algorithm assesses educational quality using *U.S. News* scores, such missing data will bias its outputs. Bias of this kind can be ameliorated by correcting missing data or switching to another more comprehensive metric. But not all bias is so easily avoided. The whole point of the hiring algorithm is that it will learn to mimic, and thus supplant, past human hiring recommendations. Insofar as *those* past human decision-makers were biased, the algorithm will learn to mimic that bias. Then, discrimination will be baked into the algorithm's fundamental perceptions of what makes a quality hire. Here, affirmative corrections of the human-induced—and machine-assimilated—bias may be necessary.⁹⁶

Big Data Affirmative Action offers a solution. With a Big Data Affirmative Action policy in place, algorithmic decision-making would have two stages, involving two separate statistical models. The first model would provide an initial decision, and a second Big Data Affirmative Action model would then adjust that initial decision to eliminate any baked-in discrimination.

⁹² See, e.g., Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan & Cass R. Sunstein, *Discrimination in the Age of Algorithms*, 10 J. LEGAL ANALYSIS 1, 3 (2018) [hereinafter Kleinberg et al., *Discrimination*]; Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig & Sendhil Mullainathan, *Human Decisions and Machine Predictions*, 133 Q.J. ECON. 237, 238 (2018) [hereinafter Kleinberg et al., *Human Decisions*].

⁹³ See, e.g., Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671 (2016); Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold & Richard Zemel, *Fairness Through Awareness*, 2012 PROC. INNOVATIONS THEORETICAL COMPUT. SCI. CONF. 214. The website *Fairness, Accountability, and Transparency in Machine Learning* maintains an archive of dozens of such studies. *Scholarship, FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY IN MACHINE LEARNING*, <https://www.fatml.org/resources/relevant-scholarship> [<https://perma.cc/EQM8-VF73>].

⁹⁴ See Salib, *supra* note 57, at 567–72; Barocas & Selbst, *supra* note 93, at 677–93; Kleinberg et al., *Discrimination*, *supra* note 92, at 7, 43.

⁹⁵ See Kleinberg et al., *Discrimination*, *supra* note 92, at 28.

⁹⁶ It might be possible to instead improve the data about post-bail reoffense. But this is difficult. Police investigation and arrest is the primary way we determine who has committed crimes. Thus, improving the data would likely entail de-biasing police's attitudes—a goal we do not currently know how to accomplish. See *supra* Section I.B.

The initial-decision algorithm could employ any statistical approach the designer liked. If she prized model interpretability over predictive accuracy, she could use a simple ordinary least squares regression. If she prized predictive accuracy over understanding the model's inner workings, she could use a cutting-edge deep neural network.⁹⁷ The model could also have as many input features as the designer liked, again allowing optimization between accuracy and interpretability.⁹⁸

The initial-decision algorithm would digest its inputs, run the numbers, and output a tentative recommendation. Then the second model—the Big Data Affirmative Action model—would step in. Big Data Affirmative Action for algorithmic decisions would work exactly the same as it would for human decisions. Analyzing the initial-decision algorithm's inputs and outputs, the Big Data Affirmative Action model would estimate the effect of race on those outputs, controlling for other “legitimate” factors. As with human decisions, the resulting coefficients of discrimination would be used to program a corrective algorithm. The corrective Big Data Affirmative Action model would then adjust the initial-decision model's recommendation, eliminating unwarranted racial penalties and de-biasing the algorithm.

Two other scholarly articles have proposed related—but distinct—statistical approaches to eliminating algorithmic discrimination.⁹⁹ These approaches share some features of Big Data Affirmative Action. But there are differences worth noting. First, the Big Data Affirmative Action approach is more flexible. It can be applied to correct discrimination by *any* kind of algorithm, including noninterpretable “black boxes” like deep neural networks. By contrast, some of the other proposals work only to correct

⁹⁷ HASTIE ET AL., *supra* note 56, at 305–16, 389–414. Generally, as the accuracy of statistical models' answers goes up, our ability to understand how they got those answers goes down. This is called the accuracy–interpretability tradeoff. The most accurate algorithms are often described as “black boxes,” having completely incomprehensible (to humans) decision functions. See Toshiki Mori & Naoshi Uchihira, *Balancing the Trade-Off Between Accuracy and Interpretability in Software Defect Prediction*, 24 EMPIRICAL SOFTWARE ENG'G 779, 780 (2019); Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, 2016 BIG DATA & SOC'Y 1, 4–7.

⁹⁸ See Aylin Alin, *Multicollinearity*, 2 WIRES COMPUTATIONAL STAT. 370, 370–71 (2010) (explaining how the addition of collinear variables reduces a model's interpretability).

⁹⁹ See Crystal S. Yang & Will Dobbie, *Equal Protection Under Algorithms: A New Statistical and Legal Framework*, 119 MICH. L. REV. 291, 346–48 (2020); see generally Devin G. Pope & Justin R. Sydnor, *Implementing Anti-Discrimination Policies in Statistical Profiling Models*, 3 AM. ECON. J. 206 (2011) (proposing a de-biasing method using a single, interpretable statistical model trained to quantify and correct the effect of race on a given decision).

decisions by highly interpretable, but less accurate, algorithms.¹⁰⁰ This advantage of Big Data Affirmative Action will become especially important as extraordinarily accurate, but radically noninterpretable, algorithms come to dominate automated decision-making.¹⁰¹

Second, the Big Data Affirmative Action approach does not require—as others do—ignoring large amounts of data when building the initial-decision algorithm.¹⁰² Such systematic exclusion of relevant data can again reduce accuracy in algorithmic decisions¹⁰³ and perhaps introduce bias of its own.¹⁰⁴

Finally, certain other proposals do not follow Big Data Affirmative Action in carefully calibrating their race-based adjustments to correct *only* discrimination by a single institution.¹⁰⁵ As discussed below, such localization is crucial for maintaining the legality of any affirmative-action-style intervention.¹⁰⁶

Big Data Affirmative Action thus represents an improvement on the extant proposals for de-biasing algorithms. Perhaps more importantly, Big Data Affirmative Action goes much further. These previous proposals are

¹⁰⁰ Professors Yang and Dobbie propose an intervention under which a single prediction model first estimates coefficients for all relevant variables, including race. Then, when the time comes to make individual decisions, each individual's actual race is replaced with a universal dummy value. Thus, the model corrects for racial discrimination by treating everyone as if they were of the same race. Yang & Dobbie, *supra* note 99, at 346–48. This approach, however, depends on the model having accurately estimated the impact of race on outcomes. Interpretable models, like the ordinary least squares model that Yang and Dobbie use, are designed to produce such accurate estimates. *See* Yang & Dobbie, *supra* note 99, at 336 n.188 (using ordinary least squares regression). But models that prioritize predictive accuracy—like neural networks, random forests, and even some regressions—do not produce such estimates. Burrell, *supra* note 97, at 5–7. Such models ignore the relative contributions of individual inputs in order to maximize the accuracy of bottom-line outputs. *Id.* Thus, neither the Yang–Dobbie nor the related Pope–Sydnor intervention will work to de-bias highly accurate, but less interpretable, models.

¹⁰¹ Relatedly, the Big Data Affirmative Action approach works even in the presence of collinearity between race and another independent variable. Under the Yang–Dobbie and Pope–Sydnor approaches, such collinearity could result in an inaccurate weight being assigned to race, and thus, an inaccurate correction for racial discrimination. *See* Alin, *supra* note 98, at 370–71. Under a Big Data Affirmative Action approach, this can be fixed by leaving the race-correlated variable in the initial decision model but dropping it from the corrective model. *See infra* Section IV.F. Yang–Dobbie and Pope–Sydnor use just one model for both decision and correction. There, dropping the race-correlated variable will again reduce the accuracy of the model's outputs.

¹⁰² Yang and Dobbie's second proposal contemplates training the decision algorithm using only data from members of the favored racial group—generally whites. *See* Yang & Dobbie, *supra* note 99, at 349.

¹⁰³ *See id.* (noting that Yang and Dobbie's intervention produces accurate predictions only “if one believes that bias . . . is not an issue among white[s] . . . and that [their data] is [therefore] . . . accurate”).

¹⁰⁴ *See* Richard A. Berk, *An Introduction to Sample Selection Bias in Sociological Data*, 48 AM. SOCIO. REV. 386, 386–88 (1983) (describing selection bias).

¹⁰⁵ *See, e.g.,* Kleinberg et al., *Human Decisions*, *supra* note 92, at 276 (proposing an algorithmic intervention that, in effect, sets a racial quota for pretrial incarceration).

¹⁰⁶ *See* City of Richmond v. J.A. Croson Co., 488 U.S. 469, 507 (1989).

limited to correcting discrimination in *algorithmic* decisions. By contrast, as described in this Article, Big Data Affirmative Action is a truly universal tool. It can be used to correct discrimination in *any* decision—human, algorithmic, or otherwise.

III. IS BIG DATA AFFIRMATIVE ACTION LEGAL?

Big Data Affirmative Action's binding constraint is not likely to be its practicability. As discussed above, across a range of contexts, the data exist and the expertise is available to design and implement effective policies.

Instead, the most likely constraints are legal. Essentially any substantial policy intervention designed to address racial, gender, religious, or related inequalities will have to grapple with constitutional and statutory antidiscrimination law. Certain anodyne interventions, like the de-biasing techniques discussed above, are unlikely to raise serious legal concerns. But unfortunately, they are also unlikely to do any good for people of color or other Americans facing discrimination. Big Data Affirmative Action is, by design, a more muscular intervention, designed to affect the world directly and to do so explicitly along racial lines. Any intervention like that is sure to raise complicated questions of legality.

This Part does two things. First, it explains the law bearing on Big Data Affirmative Action. The central puzzle here is that, despite a long history of approving affirmative action policies of all kinds, the Supreme Court has appeared in recent years to take a turn. For the past several decades, the Court has rejected *every single* affirmative action policy designed to redress discrimination that it has reviewed. True, to date the Court has maintained tentative support for affirmative action in college admissions. But as explained below, campus affirmative action is somewhat idiosyncratic—neither intended to combat discrimination nor equipped to do so. Big Data Affirmative Action and its forebears, by contrast, aim to undo discriminatory harm by implementing explicit preferences based on race or another protected characteristic.

This Part shows that the Court's apparent turn against discrimination-redressing affirmative action is just that—apparent. The governing law has not changed. Then and now, the law has demanded two things of affirmative action policies: they must be founded on a strong basis in evidence of remediable discrimination by the implementing institution and be well tailored to correcting *only* that discrimination. What changed, as shown below, were the plans under review. Older plans satisfied these criteria; newer ones have not. This Part suggests why.

After explaining the legal rules governing affirmative action, and the cases applying them, this Part shows why Big Data Affirmative Action

would pass muster. Unlike the policies that the Supreme Court has recently rejected, Big Data Affirmative Action is, by design, empirically robust and carefully tailored. Big Data Affirmative Action's statistical approach to proving and remedying discrimination does differ somewhat from the methods used by plans that the Supreme Court has approved. But insofar as it is different, Big Data Affirmative Action's approach conforms better to current legal constraints. Thus, there is every reason to believe that Big Data Affirmative Action would be legal. Indeed, the approach is so strong that the Court has already effectively endorsed it—albeit implicitly.

Before proceeding, let us pause for a brief note on methodology. The following Section is about what the law requires and why Big Data Affirmative Action satisfies it. Ardent legal realists may be skeptical that legal analysis matters at all in this arena. Perhaps the Supreme Court is simply a conservative institution whose legal opinions serve merely as pretext to reject affirmative action of all kinds.

There are three arguments against this pessimistic view. First, it does not fit very well with the Court's actual decisions. The Court could easily outlaw affirmative action, reciting colorable legal arguments, and be done with it. Yet it hasn't. Even as it has invalidated *particular* policies, it continues to vociferously reaffirm their legal permissibility in theory. Second, even if the Court does have policy views about affirmative action—and it surely does—it often exercises its prerogative to write those into the law. As discussed below, the strong basis in evidence rule reflects a small-c conservative approach to redressing discrimination. It treats discrimination as wrong and illegal but is wary of interventions that would invite rapid, uncontrolled societal change. Thus, in this case, a policy that complies with the law ought to therefore be acceptable even to conservative Justices. Finally, even if the case law in this area were a mere smokescreen to obscure the Court's pure political decisions, analyses like the one below would expose it. If, as argued here, Big Data Affirmative Action assiduously complies with the Court's stated rules, a stubborn rejection of the policy would reveal the charade.

A. *The Law of Affirmative Action*

There are facially compelling reasons for skepticism about Big Data Affirmative Action's legality. Notably, in the last thirty years, the Supreme Court has consistently struck down similarly structured policies.¹⁰⁷ That is, it has rejected affirmative action policies under which race—or another

¹⁰⁷ This trend began with *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469 (1989). See also *infra* Section III.A.3 (discussing later cases).

protected characteristic—was a decisive factor in allocating some benefit. These recent decisions stand in stark contrast to the previous thirty years, when the Supreme Court upheld many such plans.

What changed? There are two possibilities. The first—endorsed by certain dissenting Justices,¹⁰⁸ legal scholars,¹⁰⁹ and policymakers¹¹⁰—is that the Court changed its mind about the law. Under this view, the earlier decisions endorsing certain discrimination-redressing, race-based preferences should be understood as functionally overruled. If that is right, essentially all affirmative action plans of that kind—including Big Data Affirmative Action—are now forbidden. The second possibility is that the world has changed. Specifically, it has changed in some way that made it harder to design affirmative action policies that comply with the law. If this is true, then there is hope that a fresh approach—like Big Data Affirmative Action—would be upheld as legal.

The second view is the correct one. The Supreme Court has never overruled its cases endorsing certain discrimination-redressing racial preferences. On the contrary, even as it has recently rejected certain *specific* affirmative action plans, it has simultaneously insisted that they are not *generally* forbidden. The subsequent Sections trace the case law’s history, excavate a single, consistently applied legal standard, and explain why—despite that unchanging standard—the Court has lately appeared to turn against affirmative action.

1. A Note on Campus Affirmative Action

Before trying to understand the case law governing Big Data Affirmative Action, let us pause momentarily to sort out the cases that do not govern it. Today, the term “affirmative action” most readily conjures images of elite colleges, where the war over race-conscious admissions runs hot. As recently as 2016, the Supreme Court in *Fisher v. University of Texas at Austin* reaffirmed colleges’ ability to consider race in admissions.¹¹¹ Nevertheless, anti-affirmative-action activists almost immediately sued

¹⁰⁸ *Croson*, 488 U.S. at 529 (Marshall, J., dissenting) (arguing that the “decision mark[ed] a deliberate and giant step backward in this Court’s affirmative-action jurisprudence”).

¹⁰⁹ See, e.g., Anderson C. Williams, *Is Affirmative Action Dead? The Meaning of Adarand Constructors*, 19 W.J. BLACK STUD. 254, 254 (1995) (“[M]any critics of affirmative action have argued that the *Adarand* decision sounded the death knell of affirmative action.”); Kathleen M. Sullivan, *City of Richmond v. J.A. Croson Co.: The Backlash Against Affirmative Action*, 64 TUL. L. REV. 1609, 1612 (1990) (“*Croson* culminated the Court’s long-mounting trend toward limiting the justification for affirmative action . . .”).

¹¹⁰ See, e.g., Kevin Merida & Kenneth J. Cooper, *Alarmed by High Court Rulings, Black Leaders Prepare for Action*, WASH. POST, July 3, 1995, at A1; 141 CONG. REC. S8530 (daily ed. June 16, 1995) (statement of Sen. McConnell).

¹¹¹ 579 U.S. 365, 388 (2016).

Harvard University over its admissions policy.¹¹² The Court has now granted certiorari in that case, *Students for Fair Admissions, Inc. v. President and Fellows of Harvard College*,¹¹³ perhaps signaling its readiness to work a major change in the law. Indeed, the Court may now be ready to deliver on a thirty-year-old promise to eventually overrule its earlier cases and make traditional campus affirmative action illegal.¹¹⁴

But campus affirmative action—and the law governing it—is in many ways sui generis. First, the legal justification for race-based decision-making is unique. In college admissions, race-based considerations are permitted in service of the “compelling interest” of “obtaining ‘the educational benefits that flow from student body diversity.’”¹¹⁵ According to the Court, these benefits include things like “promot[ing] cross-racial understanding, . . . break[ing] down racial stereotypes, . . . promot[ing] learning outcomes, and better prepar[ing] students for an increasingly diverse workforce and society.”¹¹⁶

Big Data Affirmative Action and its forebears, by contrast, are not about promoting the socially diffuse benefits of diversity. They are about remedying the specific harms of discrimination.

This distinction—between different justifications for race-conscious decision-making—gives rise to a second important distinction. In college admissions, race may never constitute an “automatic” or “decisive” reason to admit one applicant over another.¹¹⁷ As the Court said in *Gratz v. Bollinger*, university therefore may not, for example, assign a fixed number of “points to every single applicant from an ‘underrepresented minority’ group” when point totals determine admissions decisions.¹¹⁸ Instead, race may be considered only as one among many nondecisive attributes, in a “holistic review of each applicant’s file.”¹¹⁹ Under this legal regime, college affirmative action programs are allowed only if “race is but a ‘factor of a

¹¹² See generally *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 980 F.3d 157 (1st Cir. 2020) (noting the plaintiff’s allegation that Harvard College’s undergraduate admissions process violated Title VI of the Civil Rights Act).

¹¹³ 142 S. Ct. 895 (2022).

¹¹⁴ *Grutter v. Bollinger*, 539 U.S. 306, 343 (2003).

¹¹⁵ *Fisher*, 579 U.S. at 381 (quoting *Fisher v. Univ. of Tex. at Austin*, 570 U.S. 297, 310 (2013)); see also *Grutter*, 539 U.S. at 328.

¹¹⁶ *Fisher*, 579 U.S. at 381 (quoting *Grutter*, 539 U.S. at 330).

¹¹⁷ *Gratz v. Bollinger*, 539 U.S. 244, 271–72 (2003).

¹¹⁸ *Id.* at 271.

¹¹⁹ *Grutter*, 539 U.S. at 337.

factor of a factor” in the admissions decision.¹²⁰ By contrast, under Big Data Affirmative Action—and its historical forebears, discussed below—race is *the* decisive factor for allocating benefits.

Given these important differences, the remainder of this Article largely sets aside college-based “diversity-justified” affirmative action. Instead, it focuses on “discrimination-remediating” affirmative action. That is, it uses the term “affirmative action” to refer to policies under which race is a decisive reason to allocate some benefit, with the goal of redressing discrimination. This is not to say that the two kinds of affirmative action share *nothing* in common. In both cases, the law’s conceptual structure is the same: an affirmative action policy must be aimed at achieving some legitimate end, and the policy must be well tailored to that end. But different ends authorize different means. The following Sections therefore focus on policies—and the law governing them—that share Big Data Affirmative Action’s ends.

These two varieties of affirmative action have one more thing in common: they are potential substitutes. As already noted, the Supreme Court may be poised to outlaw diversity-justified campus-based affirmative action. Suppose that is the ultimate result of the Harvard litigation. The question then becomes: What will college administrators concerned about the underrepresentation of minority students on their campuses be able to do about it?

The answer is Big Data Affirmative Action. Deprived of diversity-promoting affirmative action as an option, admissions offices will have to resort to discrimination-remediating affirmative action. Of course, the availability of discrimination-remediating affirmative action will depend on first documenting discrimination in the admissions process.¹²¹ But if such discrimination is documented, Big Data Affirmative Action is—as argued herein—the optimal approach for remedying it. Thus, if the Supreme Court

¹²⁰ *Fisher*, 579 U.S. at 375 (quoting *Fisher v. Univ. of Tex. at Austin*, 645 F. Supp. 2d 587, 608 (W.D. Tex. 2009)). This doctrine suffers from the drawback of incoherence—suggesting that race may be used in admissions only if it does not matter. The most charitable reading of *Gratz* might be that it does not flatly forbid the “decisive” use of race, nor an allocation of “points” to members of a particular racial group. Rather, the problem in that case was that the point bonus was too big. See *Gratz*, 539 U.S. at 270. But this view is in significant tension with much of *Gratz*’s language. See, e.g., *id.* at 270–71 (“[P]referring members of any one group for no reason other than race or ethnic origin is discrimination for its own sake.” (quoting *Regents of the Univ. of Cal. v. Bakke*, 438 U.S. 265, 307 (1978))).

¹²¹ This will be complicated if the prior status quo included affirmative action that benefitted groups who would otherwise be discriminated against. Thus, any new implementation of affirmative action may involve an initial fact-gathering stage, in which admissions outcomes are observed absent any affirmative action intervention.

does indeed outlaw diversity-promoting campus affirmative action, Big Data Affirmative Action will become all the more important as a policy design.¹²²

2. *The Golden Age of Affirmative Action*

In 1979, in *United Steelworkers of America v. Weber*, the Supreme Court upheld Kaiser Aluminum's contract with the United Steelworkers of America.¹²³ Under that agreement, Kaiser would reserve 50% of the seats in its craftwork training program for Black applicants, and it would hire its new craftworkers from that program. This would continue until the proportion of Black craftworkers at Kaiser's plant equaled that of the surrounding community.¹²⁴

The Court held that this plan was allowed, despite Title VII's prohibition on discrimination "because of . . . race."¹²⁵ Title VII, the Court reasoned, did not "condemn all private, voluntary, race-conscious affirmative action plans."¹²⁶ On the contrary, the statute was designed to eradicate workplace discrimination. Kaiser and the USWA's plan aimed at just that—"eliminat[ing] traditional patterns of racial segregation" in Kaiser's industry.¹²⁷

There was ample evidence of discrimination to eliminate. The Court cited a half dozen legal decisions, along with federal government reports and academic studies, documenting race-based exclusion from craftwork. In fact, the Court thought discrimination in these trades to be so well established that it was a proper subject of judicial notice.¹²⁸

Such evidence of discrimination, however, did not empower Kaiser and the USWA to impose *whatever* affirmative action plan they pleased. Their actual plan was permissible only because it was tailored to remedying the evident discrimination at hand. As the Court said, the plan would "open employment opportunities . . . which ha[d] traditionally been closed to" Black workers.¹²⁹ But it would not go beyond remedying discrimination to "unnecessarily trammel the interests of . . . white employees."¹³⁰ And the

¹²² Note, however, that since diversity-justified and discrimination-remediating affirmative action are not identical, they are not perfect substitutes. Thus, there are important questions to be answered about which groups would benefit in relative terms should admissions decisions shift from the former to the latter.

¹²³ 443 U.S. 193, 197 (1979).

¹²⁴ *Id.* at 197, 199.

¹²⁵ 42 U.S.C. § 2000e-2(a).

¹²⁶ *Weber*, 443 U.S. at 208.

¹²⁷ *Id.* at 201.

¹²⁸ *Id.* at 198 n.1.

¹²⁹ *Id.* at 208.

¹³⁰ *Id.*

Court emphasized that the plan would be in place only long enough to “eliminate [the] manifest racial imbalance.”¹³¹

Similar logic had driven the Court’s decision in *Swann v. Charlotte-Mecklenburg Board of Education*, decided eight years before *Weber*.¹³² There, the Court upheld a plan¹³³ under which Black children would be—based on their race—bused to white schools.¹³⁴ This policy was part of the United States’ years-long effort to remedy school segregation in the wake of *Brown v. Board of Education*.¹³⁵

As with *Weber*, in *Swann* there was ample evidence of discrimination to remedy. This was in some sense obvious. In Charlotte, as in many parts of the country, school segregation had once been written explicitly into law.¹³⁶ But even so, the *Swann* Court strongly emphasized the district court’s “numerous hearings” and “voluminous evidence” showing that, even after *Brown*, school segregation in Charlotte was the product of discrimination.¹³⁷

The Court emphasized this Charlotte-specific evidence because it did not believe affirmative action programs were legal when designed to address “de facto segregation.”¹³⁸ That is, such policies could not attempt to correct a mere “racial imbalance” without evidence that it was “brought about by discriminatory action.”¹³⁹ Nor would a busing plan be permitted if designed to undo “all the problems of racial prejudice.”¹⁴⁰ Only problems of prejudice in Charlotte of which there was adequate evidence could be targeted.

As with *Weber*, the *Swann* Court emphasized the affirmative action plan’s tailoring to the problem of intentional segregation in Charlotte. The

¹³¹ *Id.*

¹³² 402 U.S. 1 (1971).

¹³³ Astute observers of antidiscrimination law may object that *Swann* did not really involve an affirmative action plan. Unlike in *Weber*, *Swann*’s school busing program was not voluntarily adopted by a party to the case but was instead imposed by a district court. *Id.* at 10. That is half true. The district court did impose a mandatory desegregation scheme, but that is only because the school board failed to adopt a voluntary one. *Id.* at 6–10. The board could have adopted a voluntary one, and in fact the Supreme Court emphasized that the district had greater leeway to do so than the courts. *Id.* at 16. That is, the power of institutions to voluntarily institute race-based preferences to remedy discrimination is broader than the power of courts to impose such remedies. Nevertheless, *Swann* does blur the line between what we usually call affirmative action and what we call a judicial remedy. As discussed below, the symmetries and asymmetries between these two conceptual categories—and the gray area separating them—in fact supports Big Data Affirmative Action’s legality. See *infra* Section III.B.

¹³⁴ See *Swann*, 402 U.S. at 9.

¹³⁵ 347 U.S. 483 (1954).

¹³⁶ See *Swann*, 402 U.S. at 7.

¹³⁷ *Id.*

¹³⁸ *Id.* at 17.

¹³⁹ *Id.* at 17–18.

¹⁴⁰ *Id.* at 23.

plan set targets for racial composition within schools.¹⁴¹ But, recognizing that not *all* of the raw imbalance could be attributed to discrimination, the plan treated those targets as “a starting point . . . rather than an inflexible requirement.”¹⁴² For similar reasons, the plan called for “close scrutiny,” though not automatic condemnation, of every “one-race” school.¹⁴³

In subsequent years, the Supreme Court upheld a diverse collection of affirmative action programs across a number of domains. In *Fullilove v. Klutznick*, the Court approved a plan to reserve 10% of federal public-works funds to be spent contracting with minority-owned businesses.¹⁴⁴ Crucial to this approval were Congress’s findings that minorities’ underrepresentation in government construction work was caused by the government’s own discrimination. A congressional report found that such underrepresentation was “not the result of random chance.”¹⁴⁵ Instead, “[w]ith specific reference to Government construction contracting,” Congress found that “the practices of some agencies preclude[d]” minority participation.¹⁴⁶ The spending set-asides were also tailored specifically to address just the identified government discrimination. The plan was subject to waivers, creating a defeasible “assumption” that “‘adjustment for the effects of past discrimination’ would assure that at least 10% of the funds from the federal grant program would flow to minority businesses.”¹⁴⁷ If, however, racial disparities on certain projects were *not* caused by congressional discrimination—but instead a lack of qualified contractors—the 10% requirement could be waived. “[W]ithout this fine tuning to remedial purpose,” the Court later remarked, “the statute would not have ‘passed muster.’”¹⁴⁸

Two other cases, both decided in 1987, are also worth highlighting: *Johnson v. Transportation Agency*¹⁴⁹ and *United States v. Paradise*.¹⁵⁰ Both cases involved affirmative action employment plans for the promotion of disadvantaged minorities. The plan in *Paradise* required that 50% of the promotions in a police department go to Black officers, conditional on

¹⁴¹ See *id.* at 24.

¹⁴² *Id.* at 25.

¹⁴³ *Id.* at 25–26.

¹⁴⁴ 448 U.S. 448, 453 (1980).

¹⁴⁵ *Id.* at 465 (quoting H.R. REP. NO. 94-468, at 2 (1975)).

¹⁴⁶ *Id.* at 466 (quoting H.R. REP. NO. 94-468, at 29).

¹⁴⁷ *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469, 488–89 (1989) (quoting *Fullilove*, 448 U.S. at 487–88).

¹⁴⁸ *Id.* at 489 (quoting *Fullilove*, 448 U.S. at 487).

¹⁴⁹ 480 U.S. 616 (1987).

¹⁵⁰ 480 U.S. 149 (1987).

qualified applicants being available.¹⁵¹ *Johnson*'s plan imposed no numerical ratio but authorized race and sex as decisive reasons to promote one employee over another.¹⁵² In both cases, promotion was allowed even when Black or female employees had, for example, scored lower on competency tests than their white or male counterparts.¹⁵³

Both affirmative action plans were promulgated to redress well-documented discrimination by the implementing employers. In *Paradise*, the Court's principal opinion proclaimed, "[I]t cannot be gainsaid that white troopers promoted since 1972 were the specific beneficiaries of an official policy which systematically excluded all blacks."¹⁵⁴ The result was that, among nearly 200 upper-level officers, the department had zero Black majors, captains, lieutenants, or sergeants.¹⁵⁵ It had just four Black corporals.¹⁵⁶ In *Johnson*, likewise, the Court agreed that the transportation agency had offered "limited opportunities . . . in the past . . . for women to find employment in certain job classifications."¹⁵⁷ As a result, among 238 skilled craftworkers, none were women.¹⁵⁸

As in the earlier cases, the plans in *Johnson* and *Paradise* were tailored to address the employers' discrimination without going further.¹⁵⁹ Neither mandated gender- or race-balanced quotas in the workforce. Both made their hiring preferences conditional on the availability of qualified candidates who would be evaluated against white and male applicants along numerous dimensions.¹⁶⁰ That is, if external factors made well-qualified Black or female candidates scarce, the policy would not apply. Furthermore, both policies were designed to terminate once discrimination was eradicated from promotion decisions.¹⁶¹

These five cases—*Weber*, *Fullilove*, *Swann*, *Johnson*, and *Paradise*—are not the only ones in which the Court blessed affirmative action policies.¹⁶² But they are illustrative. As will be discussed, each has an echo in the subsequent era—a later, facially similar plan that the Court rejected. Thus, a

¹⁵¹ *Id.* at 163.

¹⁵² 480 U.S. at 622.

¹⁵³ *Id.* at 623–24; *Paradise*, 480 U.S. at 160–62.

¹⁵⁴ 480 U.S. at 170 (alteration in original) (quoting *Paradise v. Prescott*, 767 F.2d 1514, 1533 n.16 (11th Cir. 1985)).

¹⁵⁵ *Id.* at 163.

¹⁵⁶ *Id.*

¹⁵⁷ 480 U.S. at 634.

¹⁵⁸ *Id.* at 636.

¹⁵⁹ *Id.* at 637; *Paradise*, 480 U.S. at 185–86.

¹⁶⁰ *Johnson*, 480 U.S. at 636–38; *Paradise*, 480 U.S. at 177–79.

¹⁶¹ *Johnson*, 480 U.S. at 639–40; *Paradise*, 480 U.S. at 178.

¹⁶² See, e.g., *Loc. 28 of the Sheet Metal Workers' Int'l Ass'n v. EEOC*, 478 U.S. 421 (1986).

careful comparison of the Court's reasoning in each of these approval–rejection dyads can illustrate what the law requires of affirmative action policies. And it can reveal how those requirements have—or have not—changed.

3. *A Turning Tide: Wygant and the Strong Basis in Evidence Standard*

It is tempting to view *Wygant v. Jackson Board of Education* as marking a turning point against affirmative action in the case law. There, the Court invalidated a teachers' union bargaining agreement that included race-based protections against layoffs.¹⁶³

Wygant is significant not because it was the first case to invalidate an affirmative action policy. It wasn't.¹⁶⁴ But *Wygant* crystalized for the first time a single legal standard under which affirmative action policies could be evaluated. In the earlier cases discussed above, the Court's statements of legal principle ranged from hazy to evasive.

The principal opinion in *Wygant* laid down two requirements for promulgating an affirmative action policy: First, the implementing institution must “ha[ve] a *strong basis in evidence* for its conclusion that remedial action [i]s necessary.”¹⁶⁵ That evidence could not be of “societal discrimination alone,” but rather must include a “showing of prior discrimination by the [institution] involved.”¹⁶⁶ Second, the policy must be “limited and properly tailored . . . to cure the effects of [that] prior discrimination.”¹⁶⁷ That is, an institution may clean up its own discriminatory mess, but no one else's.

Wygant's strong basis in evidence rule enshrines a compromise between ambitious progressive reformism and small-c conservatism. Progressives concerned with eradicating social injustice might prefer that affirmative action be much more aggressive. Perhaps employers ought to be allowed to confer large advantages to *all* workers from historically disadvantaged groups, irrespective of who discriminated, when, and how much. However, conservatives would object to that strategy, even if they agreed that discrimination was real and wrong. Their objection might go thus: The cumulative effects of such aggressive affirmative action—*many* employers conferring *many* advantages—would go too far too fast.

¹⁶³ 476 U.S. 267, 270 (1986).

¹⁶⁴ See *Milliken v. Bradley*, 418 U.S. 717 (1974).

¹⁶⁵ 476 U.S. at 277 (emphasis added). *Wygant* produced no majority opinion. But in subsequent years, the Supreme Court adopted the strong basis in evidence standard from Justice Powell's principal opinion as authoritative. See, e.g., *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469, 510 (1989).

¹⁶⁶ 476 U.S. at 274.

¹⁶⁷ *Id.* at 281 (quoting *Fullilove v. Klutznick*, 448 U.S. 448, 484 (1980)).

Untethered from particular instances of discrimination, the net result could be a kind of “double recovery,” by which the total quantum of remedy exceeded the quantum of harm. Another concern, raised by Justice O’Connor in a different case, is that unconstrained affirmative action might invite private actors to engage in “outright racial balancing.”¹⁶⁸ Such balancing, in the conservative view, simply replaces one kind of discrimination (against historically disadvantaged groups) with another (against historically advantaged groups).

The strong basis in evidence rule, as concretized in *Wygant*, represents a middle path. It treats discrimination as real, wrong, and illegal and endorses affirmative action as a widely available, nonjudicial remedy. But it authorizes the remedy only in response to—and congruent with—evinced discrimination by the implementing institution. This blunts the double-recovery and racial-balancing objections. The result is a legal test that responds to progressive concerns but does so in a way that conservatives can, and should, accept.

Beyond crystalizing the strong basis in evidence rule, *Wygant* was also perhaps significant for a second reason. It foreshadowed an era in which the Court has consistently rejected affirmative action policies—despite having previously approved facially identical ones. The plan at issue in *Wygant* looked very much like that of *Weber*. Both afforded employment preferences to members of disadvantaged racial groups. Yet, in contrast to the *Weber* Court, the *Wygant* Court rejected the affirmative action policy under review as impermissible under its newly articulated rule.

In the intervening decades, this story repeated itself over and over. Time after time, the Court has struck down affirmative action policies that looked, at first blush, very much like ones it had previously upheld. In *City of Richmond v. J.A. Croson Co.*, the Court invalidated the city of Richmond’s law setting aside 30% of government construction spending for minority-owned businesses.¹⁶⁹ This despite the law’s facial similarity to the set-aside approved in *Fullilove*. In *Parents Involved in Community Schools v. Seattle School District No. 1*, the Court forbade Seattle’s use of race in assigning students to public schools.¹⁷⁰ Yet it had allowed just that in *Swann*. And in *Ricci v. DeStefano*, the Court disapproved a fire department’s promotion of minority firefighters over white firefighters who scored higher on a qualifying exam.¹⁷¹ Yet it had permitted both race and sex to trump test scores in *Paradise* and *Johnson*, respectively.

¹⁶⁸ *Croson*, 488 U.S. at 507.

¹⁶⁹ 488 U.S. at 477.

¹⁷⁰ 551 U.S. 701, 711–14 (2007).

¹⁷¹ 557 U.S. 557, 562 (2009).

4. *New Law or New Facts?*

What happened after *Wygant*? There are two possibilities. The first is that *Wygant* did not merely formalize a previously inchoate legal standard for evaluating affirmative action plans. Instead, it changed the law. Perhaps it raised the standard substantially, so that going forward, few plans could survive. Maybe it did something even more dramatic, overruling earlier cases and, sub silentio, outlawing affirmative action entirely.

These views have some undeniable appeal, especially in light of the Court's apparent inconsistency in upholding and striking down facially similar policies. It is no surprise then that, with mounting emphasis after each adverse Supreme Court decision, academics and other court watchers have repeatedly declared the death of affirmative action in America.¹⁷² Moreover, one thing *is* indisputable: when the Court has recently struck down affirmative action plans, it has predictably relied on *Wygant*.¹⁷³

There is, however, another explanation for the Court's apparent about-face on affirmative action. This explanation sounds not in law, but in fact. Perhaps *Wygant* really did just clarify—not change—the legal test for affirmative action. Maybe, then, it is a mere coincidence that affirmative action plans recently started failing the test despite facial similarity to plans that had passed it.

There are three reasons to think this latter view is the correct one. First, the Court has said over and over that affirmative action is still allowed. Second, the consistent-rule theory makes sense of what would otherwise be an inexplicable timeline of case law. Third—and most important—a careful reading of the cases reveals why, even under a consistent legal standard, some affirmative action plans passed muster while other similar plans failed.

Begin with what the Court keeps saying. As recently as *Ricci*—decided in 2009—the Court reaffirmed, “The [strong basis in evidence] standard leaves ample room for employers’ voluntary [antidiscrimination] efforts.”¹⁷⁴ This includes “discretion in making race-based decisions”—that is, affirmative action.¹⁷⁵ Likewise, in *Parents Involved* (decided in 2007), Justice Roberts declaimed that “no one questions that the obligation to” redress discrimination in schooling “can include race-conscious remedies—

¹⁷² See *supra* notes 108–110 and accompanying text.

¹⁷³ See, e.g., *Croson*, 488 U.S. at 494 (noting the Court's “continued adherence to the standard of review employed in *Wygant*”); *Ricci*, 557 U.S. at 582–83 (discussing the *Wygant* plurality before applying the strong basis in evidence standard).

¹⁷⁴ 557 U.S. at 583.

¹⁷⁵ *Id.*

whether or not a court had issued an order to that effect.”¹⁷⁶ According to the Court, then, neither *Wygant* nor any other case has rendered affirmative action literally or functionally illegal. It remains permissible today.

Second, the timeline. Certainly, the Court’s longest dry spell in approving affirmative action took place after *Wygant*. But not immediately after. *Johnson* and *Paradise* were both decided the year following *Wygant*. Both of those cases grappled with *Wygant*’s rejection of an affirmative action plan,¹⁷⁷ but both Courts determined that the particular plans before them were permissible. Justice O’Connor’s concurrence in *Johnson*, in fact, was devoted entirely to illustrating uniformity in the Court’s reasoning from *Wygant* back to *Weber*.¹⁷⁸ The *Johnson* majority, too, conducted itself as if the rule had not changed, acknowledging *Wygant* but then relying on *Weber* for the relevant rule.¹⁷⁹ Instead, it was not until *Croson*—decided two years after both *Paradise* and *Johnson*—that the Court kicked off its decades-long rejection streak.¹⁸⁰

Moreover, even before *Wygant*, the court had been in the business of rejecting affirmative action plans that resembled plans it had previously blessed. In *Milliken v. Bradley*, decided over a decade before *Wygant*, the Court rejected a school affirmative action program facially similar to the one it had approved in *Swann*.¹⁸¹ All of this undermines the theory that *Wygant* or the strong basis in evidence standard propagated a seismic shift in the law of affirmative action.

If the law has not changed, then the Court’s apparent decades-long crusade against affirmative action demands another explanation. The most straightforward story is that the Court believed the plans it invalidated to be factually distinct, in critical respects, from the ones it upheld. Put simply, the affirmative action plans the Court upheld—in the Court’s view—satisfied the strong basis in evidence test. And the ones it struck down did not. Careful readings of the opinions in each accepted–rejected dyad reveals precisely this dynamic.

Consider first the pairing of *Weber* and *Wygant*. *Weber*’s affirmative action plan favoring Black craftworkers came on the heels of “[j]udicial

¹⁷⁶ *Parents Involved in Cmty. Schs. v. Seattle Sch. Dist. No. 1*, 551 U.S. 701, 737 (2007). While *Parents Involved* did not produce a majority opinion, even the famously affirmative-action-skeptical Justices Scalia and Thomas joined Chief Justice Roberts’s opinion.

¹⁷⁷ *Johnson*, 480 U.S. at 626–27; *Paradise*, 480 U.S. at 182–83.

¹⁷⁸ *Johnson*, 480 U.S. at 647–57 (O’Connor, J., concurring in the judgment).

¹⁷⁹ *Id.* at 627–28 (majority opinion) (“The assessment of the legality of the Agency Plan must be guided by our decision in *Weber*.”).

¹⁸⁰ See *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469 (1989); *Parents Involved*, 551 U.S. 701; *Ricci*, 557 U.S. 557.

¹⁸¹ 418 U.S. 717, 740–45 (1974).

findings of exclusion from crafts on racial grounds . . . so numerous as to make such exclusion a proper subject for judicial notice.”¹⁸² In *Wygant*, the situation was precisely reversed. There, too, previous litigation had produced factual findings about the promulgating institution’s discrimination in hiring. “This precise issue was litigated in [two previous] suits. Both courts concluded that any statistical disparities were the result of general societal discrimination, not of prior discrimination by the [employer].”¹⁸³ Thus, on the Court’s view, the *Weber* plan was premised on evidence that the favored Black workers faced discrimination *by their employer*. The *Wygant* plan was not.

The *Fullilove–Croson* dyad turned on the same distinction. In *Fullilove*, Congress’s affirmative action plan for federal contracting rested on strong empirical findings of discrimination by the federal government. Congress began by amassing evidence of large nationwide disparities between white-owned and minority-owned businesses in securing government contracts.¹⁸⁴ But it did not stop there. Turning to the U.S. government’s own conduct, Congress identified “the practices of some agencies [that] preclude[d]” minority-owned businesses from winning federal contracts.¹⁸⁵ These included, among other things, “the exercise of discretion by government procurement officers to disfavor minority businesses.”¹⁸⁶

By contrast, in *Croson*, the City of Richmond produced *no evidence* of its own discrimination. The city identified a gross statistical disparity. But even this was weak. The disparity was not between government construction contracts with white-owned businesses and contracts with minority-owned businesses. Instead, it was between the proportion of contracts with minority-owned businesses and the proportion of the city’s residents who were minorities.¹⁸⁷ This, the Court thought, was not convincing evidence of discrimination by the city. It could just as easily be evidence of a lack of qualified minority-owned contractors in the city. That lack, of course, might also have been caused by discrimination—in education, private hiring, lending, or elsewhere. But evidence of such broad “societal discrimination alone”—as opposed to discrimination attributable to the institution implementing affirmative action—does not suffice in the strong basis calculus.¹⁸⁸ The Court found the same fault with the City’s reliance on

¹⁸² *United Steelworkers of Am. v. Weber*, 443 U.S. 193, 198 n.1 (1979).

¹⁸³ *Wygant v. Jackson Bd. of Educ.*, 476 U.S. 267, 278 (1986).

¹⁸⁴ *Fullilove v. Klutznick*, 448 U.S. 448, 465 (1980) (citing H.R. REP. NO. 94-468, at 2 (1975)).

¹⁸⁵ *Id.* at 466 (quoting H.R. REP. NO. 94-468, at 29).

¹⁸⁶ *Id.* at 467.

¹⁸⁷ *City of Richmond v. J.A. Croson Co.*, 488 U.S. 469, 501 (1989).

¹⁸⁸ *Id.* at 505.

Congress’s findings, from *Fullilove*, of nationwide discrimination by governments against minority-owned contractors.¹⁸⁹ “In sum, none of the evidence presented by the city point[ed] to any identified discrimination *in the Richmond construction industry*.”¹⁹⁰

Furthermore, the *Croson* plan was poorly tailored compared with the plan in *Fullilove*. Recall that in *Fullilove*, the 10% spending set-aside was flexible. It could be waived when too few minority-owned businesses were available or when such businesses’ high prices were not attributable to the government’s discrimination.¹⁹¹ Such flexibility helped ensure that the *Fullilove* plan was sensitive to legitimate differences between contractors aside from race.¹⁹² *Croson*’s plan was much more rigid. Its 30% quota was not tied to any careful estimate of how many minority-owned contractors would be available absent Richmond’s discrimination.¹⁹³ Nor was there flexibility to adjust the set-aside as the city learned more about contractor availability under a nondiscriminatory bid system. “No partial or complete waiver” was available “other than in exceptional circumstances.”¹⁹⁴ The Court thus saw Richmond’s plan not as a careful remedy for the city’s own discrimination, but as an exercise in “outright racial balancing.”¹⁹⁵

The Court’s other opposing pairs of affirmative action cases follow the same pattern. In *Swann*, the City of Charlotte relied on city-specific evidence of intentional school segregation.¹⁹⁶ In contrast, in *Parents Involved*, one of two districts “ha[d] not shown that [it was] ever segregated by law.”¹⁹⁷ The other had been segregated, had been subject to a desegregation order, and a court found that it “ha[d] eliminated the vestiges of its prior segregated school system.”¹⁹⁸

The affirmative action policies in *Paradise* and *Johnson* were premised on evidence that Black–white and female–male workplace disparities were caused by discriminatory employment policies.¹⁹⁹ As such, the employers were entitled to implement explicit preferences for promoting Black and female employees, even over similarly situated candidates with modestly higher test scores.

¹⁸⁹ *Id.* at 504.

¹⁹⁰ *Id.* at 505 (emphasis added).

¹⁹¹ *Id.* at 489.

¹⁹² *Id.* at 508.

¹⁹³ *See id.* at 504–05.

¹⁹⁴ *Id.* at 478.

¹⁹⁵ *Id.* at 507.

¹⁹⁶ *See supra* notes 136–140 and accompanying text.

¹⁹⁷ *Parents Involved in Cmty. Schs. v. Seattle Sch. Dist. No. 1*, 551 U.S. 701, 720 (2007).

¹⁹⁸ *Id.* at 732.

¹⁹⁹ *See supra* notes 149–160 and accompanying text.

The fire department in *Ricci*, too, wished to promote minority officers over similar white officers with higher test scores.²⁰⁰ But this was not because of a history of discrimination by the department. It was instead because of a worry that the *test itself* was discriminatory. The department's promotion exam did produce statistically disparate results.²⁰¹ But, the Court ruled, there was overwhelming evidence that the fire department had carefully calibrated its test to assess important job-related knowledge.²⁰² Thus, the Court held, the test's disparate impact was, at most, evidence of a limited supply of qualified Black and Brown officers.²⁰³ Lacking evidence that the department was remedying its own discrimination, the Court viewed the policy as an illegal program which could be used to "obtain[] the employer's preferred racial balance."²⁰⁴

In sum, although it has been decades since the Supreme Court has seen an affirmative action policy that it liked, affirmative action has not been outlawed. Instead, when the Court has recently struck down plans that resembled ones it had previously upheld, it has done so because the plans themselves were, in fact, factually different. The permissible plans were based on strong evidence that the institution promulgating the plan discriminated against the plan's favored group. These plans, further, were tailored to that evidence, implementing remedies sufficient to overcome the evinced discrimination, but not more. By contrast, the affirmative action policies that the Court has struck down were premised on weak empirical evidence of discrimination. That evidence often suggested, at most, that someone other than the implementing institution discriminated against a disfavored group. And of course, plans lacking concrete evidence of discrimination by the implementing institution cannot, by definition, be tailored to remedying just that discrimination. Where the quantum of such discrimination is unknown—or zero—no tailoring is possible.

5. *The Era of Quiet(er) Racism*

The above analysis of a half century of case law provides only a partial answer to the central mystery of the Court's apparent turn against affirmative action. As argued above, the problem with the plans that the Court struck down in recent years was that they were not founded on sufficient empirical evidence of discrimination. But *why not?* Why has it apparently become

²⁰⁰ *Ricci v. DeStefano*, 557 U.S. 557, 574 (2009).

²⁰¹ *Id.* at 586–87.

²⁰² *Id.* at 587–89.

²⁰³ *See id.*

²⁰⁴ *Id.* at 582.

harder in the past 30 years to assemble a strong basis in evidence on which an affirmative action plan may be based?

Perhaps the world has changed, making the facts more difficult to unearth. The early affirmative action cases—*Swann*, *Weber*, *Fullilove*, etc.—were decided during an era of extraordinarily overt racism and equally overt discrimination. Consider one 1969 case of which the *Weber* Court took judicial notice, in which the defendant union conceded in its brief that “[a]ll [Black applicants] were denied referral [for membership] admittedly because they were negroes.”²⁰⁵ A union official in another such case admitted that “I have 125 white men sitting here who have been paying dues in this union for years and I cannot send out the black ones before . . . them.”²⁰⁶ The prejudice was out in the open. There was simply no subterfuge.

Or consider schools. *Swann* was one of a long line of judicial decisions attempting to implement the promise of *Brown*. In *Brown*, of course, the Supreme Court ordered desegregation “with all deliberate speed.”²⁰⁷ This did not happen. Instead, people took to the streets in protest, armed mobs tried to block Black students from entering white schools, and 101 congressmen signed *The Southern Manifesto*, denouncing *Brown*.²⁰⁸

Racism and discrimination do not look like that today—at least not as often.²⁰⁹ Since the middle of the twentieth century, Americans have steadily become less willing to openly espouse racist or discriminatory views. According to the General Social Survey (GSS), between the 1970s and 1990s, public support for whites’ right to segregate neighborhoods fell by around two-thirds.²¹⁰ Support for laws prohibiting interracial marriage followed the same trend. In 2008, about 33% fewer people believed that homeowners should have the right to discriminate in selling their houses than believed so in 1972.²¹¹ And by the mid-1980s, so few people supported school segregation that the GSS stopped asking.²¹²

²⁰⁵ *United Steelworkers of Am. v. Weber*, 443 U.S. 193, 198 n.1 (1979); *Loc. 53 of the Int’l Ass’n of Heat & Frost Insulators v. Vogler*, 407 F.2d 1047, 1050 n.6 (5th Cir. 1969). This case notably arose in the same district court as *Weber*.

²⁰⁶ *United States v. United Bhd. of Carpenters of Am.*, Loc. 169, 457 F.2d 210, 215 (7th Cir. 1972).

²⁰⁷ *Brown v. Bd. of Educ.*, 349 U.S. 294, 301 (1955).

²⁰⁸ See *Resistance to School Desegregation*, EQUAL JUSTICE INITIATIVE (Mar. 1, 2014), <https://eji.org/news/history-racial-injustice-resistance-to-school-desegregation/> [<https://perma.cc/B6MZ-QGP5>].

²⁰⁹ There are, of course, exceptions. See, e.g., Fausset, *supra* note 1 (describing white-supremacist rallies).

²¹⁰ Lawrence D. Bobo, Camille Z. Charles, Maria Krysan & Alicia D. Simmons, *The Real Record on Racial Attitudes*, in *SOCIAL TRENDS IN AMERICAN LIFE: FINDINGS FROM THE GENERAL SOCIAL SURVEY SINCE 1972*, at 47 (Peter V. Marsden ed., 2012).

²¹¹ *Id.*

²¹² *Id.*

These findings do not mean that America today is post-racial or post-racist. On the contrary, the same results include double-digit percentages of Americans supporting every surveyed prejudiced position except school segregation.²¹³ And as discussed above, actual discrimination—in employment, healthcare, criminal justice, and more—continues to thrive.

In fact, the GSS reports do not even prove that Americans are *any* less racist than they were fifty years ago.²¹⁴ They simply show that people are less willing to *say* that they are racist. This comports with everyday experience. Today, it is difficult to imagine 101 sitting congresspeople signing the *Southern Manifesto*. But power brokers still enact intentionally discriminatory policies. Consider, for example, the “Muslim ban.”²¹⁵ It is just that, today, policymakers feel compelled to use dog whistles to obfuscate discriminatory motives that they historically might have espoused openly. Even Donald Trump—not known for rhetorical circumspection—felt compelled to dress the “Muslim ban” in the facially race- and religion-neutral language of national security.²¹⁶

Here, then, lies the crux of the affirmative action policymaker’s dilemma, circa 2022. Racial disparities remain pervasive. Thoughtful policymakers believe these disparities are driven by discrimination. But no one will *admit* discrimination anymore. So, as compared with the past, discrimination is now harder to prove. And as a result, the strong basis in evidence required to promulgate an affirmative action policy has become more difficult to assemble.

This model can explain the apparent disjunction between cases like *Weber*, *Swann*, and *Fullilove* and cases like *Wygant*, *Parents Involved*, and *Croson*. In *all* of these cases, policymakers presented substantial evidence of gross statistical disparities between advantaged and disadvantaged groups. But raw disparities, on their own, never constitute a strong basis in evidence. As discussed above, such raw disparities can *either* be caused by an implementing institution’s discrimination or by other factors like broad societal discrimination.

In the early years, casual, overt, and often de jure racism filled this evidentiary gap. Then, policymakers had evidence of *both* statistical disparities *and* overt discriminatory acts designed to produce them.

²¹³ *Id.*

²¹⁴ They might be. Or maybe not. Self-report is, at best, a noisy proxy for socially and legally forbidden attitudes.

²¹⁵ See Declan Walsh, *Barred from U.S. Under Trump, Muslims Exult in Biden’s Open Door*, N.Y. TIMES (Jan. 23, 2021, 8:50 PM), <https://www.nytimes.com/2021/01/23/world/middleeast/trump-muslim-ban-biden.html> [<https://perma.cc/P4B6-GXUD>].

²¹⁶ *Id.* (noting that the policy’s “stated aim was to keep terrorists out”).

Together, these provided a strong basis in evidence of localized discriminatory harm.

In the later years, however, institutions promulgating affirmative action policies lacked the evidence of overt discriminatory intent. Lacking such evidence, would-be affirmative action policymakers are again stuck *only* with gross statistical disparities. Without more, the resulting policies failed the strong basis in evidence test and were struck down.

B. *Why Big Data Affirmative Action Passes Muster*

As discussed above, the strong basis in evidence rule, as articulated in *Wygant*, requires two things:²¹⁷ First, affirmative action must be based on evidence that there is in fact a discriminatory harm to be remedied.²¹⁸ That harm must flow from the institution promulgating the plan—as opposed to another institution or society as a whole.²¹⁹ Second, the plan must be “limited and properly tailored . . . to cure the effects of [that] prior discrimination.”²²⁰ That is, the plan’s explicit race-based *preferences* must not be substantially larger than necessary to counteract the institution’s documented race-based *penalties*.

Big Data Affirmative Action policies would, by their very design, do both. Indeed, they would not merely satisfy these dual requirements; they would do so substantially *better* than the affirmative action policies that the Supreme Court has previously blessed.

Begin with the first requirement—evidence of discrimination by the implementing institution. As described above, Big Data Affirmative Action’s approach here would generally be statistical.²²¹ Policy designers would rely on rich datasets about the people—employees, criminal

²¹⁷ The Supreme Court has occasionally found itself confused as to how many tests there are for evaluating affirmative action. Is it just one—the strong basis in evidence rule—or are there two standards: one statutory and one constitutional? *See, e.g., Ricci v. DeStefano*, 557 U.S. 557, 563 (2009) (purporting to decide a statutory, but not a constitutional, question). There is just one. For example, both *Ricci* (a statutory case) and *Crosby* (a constitutional case) applied the same strong basis in evidence test. True, in constitutional cases, the Court sometimes speaks in terms of transsubstantive tiers of scrutiny. But rightly understood, the strong basis in evidence standard is just a concrete way of cashing out strict scrutiny in the affirmative action context. *Cf. Peter N. Salib, The Pigouvian Constitution*, 88 U. CHI. L. REV. 1081, 1111–13 (2021) (describing other concretizations of tiered scrutiny tests).

²¹⁸ *Wygant v. Jackson Bd. of Educ.*, 476 U.S. 267, 277 (1986).

²¹⁹ *Id.* at 274.

²²⁰ *Id.* at 281 (quoting *Fullilove v. Klutznick*, 448 U.S. 448, 484 (1980)).

²²¹ Randomized experimental approaches might also be appropriate for some applications. For example, a company wishing to eliminate discrimination in interview requests based on written applications could follow the designs of the resume experiments discussed above. *See supra* note 85 and accompanying text.

defendants, students, or others—to whom the policy would apply.²²² These datasets would catalog the attributes of those people that might have factored into the target outcomes—hiring, salary, sentence, admission. Such cataloged attributes would include both “legitimate” drivers of disparities—experience, criminal history, SAT scores—and illegitimate drivers of disparities—race, gender, religion. Statistical analysis—like regression models—would then be used to determine which inputs drove outcomes, and to what extent.

This kind of modeling draws precisely the distinction that the law requires. The statistical controls hold constant all of the most plausible “legitimate” factors—that is, factors *other than* the implementing institution’s discrimination. Differences driven by such factors are taken as given at the locus of an employer, court, or hospital implementing an affirmative action plan. And once these plausible “legitimate” explanations are accounted for, if a race-correlated gap remains, this is strong evidence that discrimination *by the institution under examination* is its cause.

How does this statistical approach to evincing discrimination compare with the evidence that the Court has historically approved? The statistical approach is fundamentally eliminative, ruling out plausible stories until only one remains. Most of the approved evidence in the case law, however, has been *additive*. On top of gross racial disparities, policymakers offered *additional* evidence of racial animus. The presence of some discriminatory intent supported an inference that gross racial disparities were actually caused by the institution’s discrimination.

Which approach produces stronger empirical evidence of an institution’s discrimination? Admittedly, the kind of statistical modelling underpinning Big Data Affirmative Action cannot eliminate every conceivable nondiscriminatory explanation for a given racial disparity. It is always logically possible—if often improbable—that, even after controlling for every “legitimate” factor one can imagine, the residual disparities had some other, yet-unimagined cause. There will always be unknown unknowns.

The direct evidence of discriminatory intent that the Supreme Court approved in cases like *Weber*, *Swann*, *Fullilove*, *Johnson*, and *Paradise* may, at first glance, seem stronger. Such evidence counts for something, to be sure. But it is not everything. And it suffers from its own serious epistemic limitations.

First, in these cases, there was often an evidentiary mismatch between the alleged discriminator and evidence of discrimination. Consider again the

²²² See *supra* Part II.

open admissions of racial animus of which the *Weber* Court took judicial notice. These admissions were from previous court cases, involving officials from unions *other* than the United Steelworkers.²²³ Those unions, in fact, had nothing to do with the case. True, all were craftwork unions. And true, some of the cited conduct arose, like in *Weber*, in Louisiana.²²⁴ And true, the evidence of explicit racism in many white Southern unions was voluminous. But none of this amounts to a smoking-gun admission that Kaiser or the local United Steelworkers *themselves* acted with discriminatory intent. Indeed, the mere fact that *this* union, among all unions in mid-century Louisiana, wished to implement affirmative action suggests the opposite.

Even if the United Steelworkers *did* intend to discriminate, it would not logically follow that they succeeded. Perhaps in *Weber*, as in other affirmative action cases, there was a general shortage of qualified Black workers in the craftwork industry.²²⁵ This is eminently plausible, since the widespread, judicially noted Southern racism in craftwork unions would likely have deterred Black workers from that career path. Under such conditions, even a *completely nondiscriminatory* union would hire out many more white workers than Black ones. Perhaps this dynamic explained some or all of the gross racial disparities identified in *Weber*.

Thus, the older, facially more straightforward approach to establishing a strong basis in evidence of discrimination produces its own uncertainty. *Weber* and related cases therefore show that absolute certainty is not required. A strong basis in evidence consists of defensible—if imperfect—empirical inferences.

Big Data Affirmative Action likewise rests on such inferences. It is always possible, in theory, that some “legitimate” factor not accounted for in a model—rather than discrimination—explains some racial disparity. But when in sentencing, for example, actual criminal conduct, a defendant’s criminal history, the presence of multiple defendants, differences in attorneys, education, age, geography, income, and employment cannot explain racial disparities,²²⁶ how likely does an alternate “legitimate” explanation seem? Surely not more likely than the possibility that widespread Southern racism against Black craftworkers, as opposed to discrimination by a *particular* employer and union, caused a racial disparity in hiring.

This point can be made even more strongly: Statistical evidence is simply *evidence*. For proof that the law of affirmative action accepts this

²²³ See *supra* notes 205–206 and accompanying text.

²²⁴ See *supra* note 205.

²²⁵ *Ricci v. DeStefano*, 557 U.S. 557, 587–88 (2009).

²²⁶ See *supra* notes 63–75 and accompanying text.

equivalence, one needs to look no further than to its sibling, civil antidiscrimination law. In many antidiscrimination lawsuits, a plaintiff may win by providing proof a policy's statistical "disparate impact."²²⁷ The statistical models used there are almost identical to the models that would underpin Big Data Affirmative Action. They proceed by showing that, controlling for plausible "legitimate" factors, a given policy produces unequal results along protected lines.²²⁸

If there were ever any doubt that this style of statistical proof could adequately support affirmative action, the *Ricci* Court put it to rest. There, the Court held, employers have "discretion in making race-based decisions" when "there is a strong basis in evidence of disparate-impact liability."²²⁹ That is, just as statistical evidence of discrimination can support civil liability, so too can it authorize affirmative action. Given all of this, there can be little doubt that Big Data Affirmative Action's statistical models can satisfy the empirical element of the strong basis in evidence rule.

What about the second element of a strong basis in evidence—tailoring? Here, Big Data Affirmative Action quite clearly fares better than even the best policies that the Court has previously approved. In early cases, the tailoring that the Court blessed was embarrassingly shoddy. In *Weber*, for example, the Court approved a set-aside of 50% of training opportunities, to be maintained until the proportion of Black craftworkers and local Black population were equalized.²³⁰ The 50% figure was apparently arbitrary. It certainly was not based on any prediction of what the discrimination-free labor composition would have been. On the contrary, the population-equalization metric implicitly attributed the *entire* Black–white craftworker disparity at Kaiser to Kaiser's discrimination. It is difficult to imagine, however, that other factors—like pervasive racism in other unions or mid-century Louisiana more broadly—did not cause at least *some* of the Black–white disparity at Kaiser.

Other policies, like the one in *Fullilove*, were more sophisticated. Recall that *Fullilove*'s 10% spending set-aside for minority-owned contractors was subject to waivers.²³¹ If, for a given project, there simply were not enough of such contractors who could do the needed work, the funds needed not be spent.²³² The same rule applied if certain minority-owned contractors charged

²²⁷ See, e.g., *Tex. Dep't of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 545–46 (2015) (affirming the cognizability of disparate impact claims under the Fair Housing Act).

²²⁸ See Salib, *supra* note 57, at 528–29.

²²⁹ 557 U.S. at 583.

²³⁰ See *United Steelworkers of Am. v. Weber*, 443 U.S. 193, 197 (1979).

²³¹ *Fullilove v. Klutznick*, 448 U.S. 448, 487–88 (1980).

²³² *Id.* at 488.

extreme prices not attributable to past discrimination.²³³ Via these waivers, the *Fullilove* plan tailored itself, at least roughly, to remedying just the effects federal government's own discrimination in contracting.

Even this is pretty clunky. How, for example, could government officials determine that high prices were or were not caused by past discrimination? They were certainly not performing complex scholarly studies of each contractor's history and cost structure. Likely, officials would simply listen to a given contractor's story and trust their gut. This approach is imprecise and thus likely to produce an affirmative-action remedy either substantially too big or too small to address the targeted discrimination.

Big Data Affirmative Action policies would be much better tailored. That is because Big Data Affirmative Action's statistical models measure both the *presence* and the *quantum* of discrimination. Knowing the quantum of discriminatory harm, Big Data Affirmative Action policymakers can easily tailor their remedies. If a court discriminatorily imposes an additional seven months' incarceration on Black defendants, then the Big Data Affirmative Action adjustment is simply seven fewer months. If Google discriminates in salary against Hispanic engineers to the tune of \$10,000 per year, then the Big Data Affirmative Action adjustment is \$10,000 more. And so on. The measured harm and the remedy are precise inverses of one another. No ex-post fact-finding or gut-trusting is needed.

Thus, under both prongs of the strong basis in evidence standard, Big Data Affirmative Action performs at least as well as—and often much better than—policies that the Court has already approved. Big Data Affirmative Action is therefore not merely an evolution of, but an *improvement on*, old-fashioned affirmative action along every legally operative dimension.

C. A Goldilocks Problem

Before moving to normative puzzles for Big Data Affirmative Action, let us pause to consider a yet-unresolved legal one: The law requires that affirmative action plans be based on “strong” evidence that the promulgating institution discriminated. But how strong, exactly? Strong evidence of discrimination, in addition to enabling affirmative action, can of course *also* trigger civil liability under antidiscrimination laws like Title VII.

This seems like a tricky situation. Institutions wishing to implement affirmative action policies—including the Big Data variety—must, on the one hand, compile evidence that they discriminated. Otherwise, their

²³³ *Id.* To see how discrimination could raise prices for some contractors, consider a white contractor who, through preferential treatment, secured enormous amounts of government work. This could lead to economies of scale and lower costs for that company, resulting in comparatively higher prices for minority-owned businesses.

affirmative action policies will *themselves* constitute illegal discrimination. On the other hand, if the evidence of discrimination is too convincing, it seems like that too will raise the specter of liability. The trick then, as with Goldilocks and the bears, is to compile evidence of one's own discrimination that is neither too hot nor too cold, but just right.

How hard is it to navigate this dilemma? That depends on the space between the two operative legal standards. If there were no daylight between the evidence constituting a strong basis and the evidence triggering liability for antidiscrimination violations, the balancing act would be difficult, indeed. Then, no right-thinking institution would ever undertake to correct its own discrimination via affirmative action.

This would be a bad situation. Happily, it is not the situation we have—at least not when Big Data Affirmative Action is on the table. First, not all antidiscrimination laws could even *potentially* threaten promulgators of Big Data Affirmative Action plans with liability. Certain statutes, like Title VII, do allow plaintiffs to prove their claims via statistical proof—the disparate impact approach.²³⁴ Here, the statistical models undergirding Big Data Affirmative Action could, at least in theory, be repurposed as weapons against the policies' designers. But under other antidiscrimination laws, like the Equal Protection Clause, disparate impact proof is not sufficient on its own to prove discrimination.²³⁵

Nevertheless, many statutes allow statistical proof, and those statutes have a wide reach.²³⁶ Here, the Supreme Court is aware of the potential Goldilocks problem. It has insisted assiduously that the gap between affirmative action's minimum evidence rule and statutory thresholds for liability is wide enough to be navigable. Title VII is again the exemplar. Of that statute, the *Weber* Court wrote, "The very statutory words intended as a spur . . . to cause 'employers and unions to self-examine . . . their employment practices and to endeavor to eliminate . . .' [discrimination] cannot be interpreted as an absolute prohibition against all private, voluntary, race-conscious affirmative action."²³⁷ The Court said much the same in *Johnson*, where it wrote that "[a] corporation concerned with maximizing

²³⁴ See *McClain v. Lufkin Indus., Inc.*, 519 F.3d 264, 279–80 (5th Cir. 2008).

²³⁵ See *Washington v. Davis*, 426 U.S. 229, 242 (1976).

²³⁶ See, e.g., *Tex. Dep't of Hous. & Cmty. Affs. v. Inclusive Cmty. Project, Inc.*, 576 U.S. 519, 545–46 (2015) (fair housing); *Wal-Mart Stores, Inc. v. Dukes*, 564 U.S. 338, 350 (2011) (employment); *Ramirez v. GreenPoint Mortg. Funding, Inc.*, 633 F. Supp. 2d 922, 926–27 (N.D. Cal. 2008) (lending). Governments, which are by default regulated via the Equal Protection Clause, may also be subject to such statutes when they are acting within the statute's scope. See, e.g., 42 U.S.C. § 2000e(a) (regulating government employers).

²³⁷ *United Steelworkers of Am. v. Weber*, 443 U.S. 193, 204 (1979) (quoting *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 418 (1975)).

return on investment . . . is hardly likely to adopt a[n] [affirmative action] plan if in order to do so it must compile evidence that could be used to subject it to a colorable Title VII suit.”²³⁸ Thus, the evidence necessary to support an affirmative action plan “need not be such that it would support a *prima facie* [discrimination] case against the employer.”²³⁹ And most recently, in *Ricci*, the Court reiterated that the strong basis in evidence standard “limits . . . discretion [to promulgate affirmative action policies] to cases in which there is a strong basis in evidence of . . . liability, but it is not so restrictive that it allows [institutions] to act only when there is a provable, actual violation.”²⁴⁰

So, the Court maintains that there is breathing room between the evidence necessary to promulgate an affirmative action plan and that which would support a disparate impact claim. But what, then, is the difference between the two? And how would these differences matter when evaluating Big Data Affirmative Action policies?

The answer, in part, is that fancy statistics alone cannot win a discrimination suit. Even under laws that allow disparate impact proof of discrimination, the numbers, without more, are not enough. In *Wal-Mart v. Dukes*,²⁴¹ the plaintiffs presented disparate impact evidence of employment discrimination very much like the evidence that might support a Big Data Affirmative Action plan. They assembled a regression model showing “disparities between men and women at Wal-Mart . . . [that] c[ould] be explained only by gender discrimination.”²⁴²

Yet the *Wal-Mart* Court rejected the employees’ Title VII claims.²⁴³ This was not because the Court discredited their statistical analysis.²⁴⁴ It was instead because the plaintiffs had failed to satisfy *another* requisite element of a disparate impact claim. They had not identified a single, particularized “pattern or practice” or “corporate policy” that caused the disparate outcomes.²⁴⁵ To succeed, a disparate impact claim brought by multiple plaintiffs requires some narrative “glue holding the alleged *reasons* for all those [allegedly discriminatory] decisions together.”²⁴⁶

²³⁸ *Johnson v. Transp. Agency*, 480 U.S. 616, 633 (1987).

²³⁹ *Id.* at 632 (emphasis added).

²⁴⁰ *Ricci v. DeStefano*, 557 U.S. 557, 583 (2009).

²⁴¹ 546 U.S. 338 (2011).

²⁴² *Id.* at 356.

²⁴³ Formally, the Court denied class certification. But it did so by looking to the merits and holding that the class could not carry its substantive burden of identifying a single discriminatory policy or practice. *Id.* at 352.

²⁴⁴ *Id.* at 356 (“Even if they are taken at face value, these studies are insufficient . . .”).

²⁴⁵ *Id.* at 352, 355.

²⁴⁶ *Id.* at 352.

Plaintiffs may carry this burden by showing, for example, that a single “biased testing procedure” caused the identified statistical disparities.²⁴⁷ They may even present evidence that the relevant decision-makers, having some discretion, all exercised that discretion to discriminate *in the same way*.²⁴⁸ This may be shown through substantial testimony of overt and related discriminatory acts against a large proportion of the plaintiffs.²⁴⁹ What is *not* enough is what the plaintiffs in *Wal-Mart* presented: a story of many different people discriminating at many different times, each in their own ways.²⁵⁰

Here, we find a fortuitous symmetry. Title VII will *not* impose liability to remedy discrimination emerging from chaotic, multifactorial processes. But such situations are precisely where Big Data Affirmative Action would be the *most useful*. When a company has a single policy driving all of its discriminatory disparities, Big Data Affirmative Action is unnecessary: The company can just get rid of the policy.²⁵¹ On the other hand, in contexts where innumerable actors are making innumerable, unrelated, and biased decisions, a Big Data Affirmative Action policy could well be the *only* way to remedy the resulting harm. Thus, Big Data Affirmative Action is *most* likely to raise the threat of legal liability when it is *least* needed. And conversely, in the many, many contexts where Big Data Affirmative Action is sorely needed, the threat of liability is minimal.

Moreover, even in contexts where a disparate impact suit might be viable, implementing Big Data Affirmative Action should *decrease* one’s net expected legal liability. Suppose that a statistical study underpinning a company’s Big Data Affirmative Action policy showed that it discriminated against women by paying them 10% less than men. And suppose that the women sued to recover for that discrimination. Suppose further that they somehow overcame the *Wal-Mart* problem and won. Then, their presumptive damages would just be the amount they were underpaid because of their race, sex, or other protected characteristic. By hypothesis, however,

²⁴⁷ *Id.* at 353 (quoting *Gen. Tel. Co. of the Sw. v. Falcon*, 457 U.S. 147, 159 n.15 (1982)).

²⁴⁸ *Id.*

²⁴⁹ *Id.* at 358.

²⁵⁰ *Id.* at 356, 358. The “policy” or “practice” requirement in disparate impact suits is likely rooted in principles of procedural justice requiring high accuracy in adjudication at the level of individual claims. See Salib, *supra* note 57, at 536.

²⁵¹ Note that the employer in *Ricci* was unable to scrap its test precisely because it had spent so much energy developing a test under which any racial disparity would be *warranted* from the perspective of employment. That is, it was careful to craft a test that evaluated necessary job skills, such that, in the Court’s view, the resulting Black–white divide was attributable not to the *department’s* discrimination, but to broader societal inequality. See *supra* notes 200–204 and accompanying text. Thus, the *Ricci* problem would not apply to an actual discriminatory policy.

the institution's Big Data Affirmative Action policy would *already* have adjusted women's salaries *upward* to counteract the 10% penalty. Thus, once a Big Data Affirmative action policy is in place, any discrimination it identifies is *already* redressed, and there is no longer a lawsuit to bring. Not only would the plaintiffs lack damages, but without any lingering injury, they might even lack Article III standing to sue.²⁵²

Institutions considering implementing Big Data Affirmative Action should therefore rest easy. The evidence needed to implement the policy is different from the evidence that would invite civil suits. Moreover, the mere implementation of a Big Data Affirmative Action policy acts as a prophylactic against antidiscrimination suits. The policies thus *reduce*, not *increase*, implementing institutions' expected legal liability.

IV. NORMATIVE OBJECTIONS

This Part explores potential normative challenges to the implementation of Big Data Affirmative Action. Some of these are old—aired regularly against affirmative action of all styles. For example, critics like Justice Clarence Thomas have long argued that affirmative action can harm its intended beneficiaries by signaling that their accomplishments are unearned.²⁵³ Others have long argued that affirmative action is simply a kind of “reverse” racism, and thus equally immoral to the wrong it is designed to remedy.²⁵⁴ This Part argues that, whether or not these critiques carry water against more traditional affirmative action—and they may not—they have no force against the Big Data variety.

Other potential challenges are newer and more specific to Big Data Affirmative Action. One is a normative cousin of the reverse racism question: Does Big Data Affirmative Action adequately match remedies to harms? The answer is yes, and it does so according to settled positive legal principles. Other critiques sound in effectiveness: Can Big Data Affirmative Action really do what it promises? And even if it can, is that enough to make a serious difference in American racial inequality? The answer to both questions is yes, so long as Big Data Affirmative Action is adopted widely and implemented carefully.

²⁵² See *Lujan v. Defs. of Wildlife*, 504 U.S. 555, 564 (1992).

²⁵³ See *Fisher v. Univ. of Tex. at Austin*, 570 U.S. 297, 333 (2013) (Thomas, J., dissenting) (arguing that affirmative action “stamps blacks and Hispanics with a badge of inferiority” (quoting *Adarand Constructors, Inc. v. Peña*, 515 U.S. 200, 215 (1995))); Ariane De Vogue, ‘*Silent Justice Outspoken on Affirmative Action*, ABC NEWS (Sept. 30, 2007), <https://abcnews.go.com/TheLaw/story?id=3667079> [<https://perma.cc/8J6R-RBNZ>] (detailing Justice Thomas’s lament that his Yale law degree “bore the taint of racial preference”).

²⁵⁴ See AMY E. ANSELL, RACE AND ETHNICITY: THE KEY CONCEPTS 136 (2013).

The final challenge is somewhat deeper. An implicit premise of the Big Data Affirmative Action proposal is that people cannot change very much. Discriminatory attitudes, whether conscious or unconscious, are sticky, and we do not know yet how to undo them. Big Data Affirmative Action essentially gives up on the project of perfecting human decision-makers and instead opts for algorithmic intervention. Humans in such a system become moral cyborgs, reliant in the long run on mechanical enhancements to their ethical decision-making. Should we be comfortable with this acceptance of our own ultimate imperfectability? And if so, does the law sanction such thinking? The answer to the former question is maybe, depending on one's goals and normative commitments. The answer to the latter is yes.

A. *“Reverse Racism” and Signaling Effects*

Perennial critics of old-fashioned affirmative action unfailingly raise three related arguments against it. The first argument is that affirmative action harms, instead of helps, its intended beneficiaries. Affirmative action does this, critics like Justice Thomas argue, by inviting doubt that the beneficiaries of affirmative action deserve their achievements. Consider a young Black attorney who is hired at an elite white-shoe firm. If the firm practices affirmative action, perhaps her colleagues will assume that she was hired because of her race and think less of her intelligence. This might be a very bad outcome for her. After all, one important reason to join such a firm is to gain the respect and admiration of lawyerly elites.

The second traditional critique of affirmative action focuses on the people who do *not* benefit from it. If an affirmative action policy benefits members of one racial group, the argument goes, it must necessarily harm members of the other groups.²⁵⁵ This argument has the most appeal in zero-sum settings, potentially including elite legal hiring. If the firm has a fixed budget for new associates, every person who is hired displaces someone else. On the other hand, this argument has less force in non-zero-sum settings like criminal sentencing. Lowering one defendant's sentence does not increase anyone else's. However, antidiscrimination law does recognize some claims based on unequal treatment even in non-zero-sum contexts.²⁵⁶

The third version of this critique takes a more abstract view. On this account, affirmative action—even when designed to redress racism—actually *constitutes* racism because it differentiates treatment on account of

²⁵⁵ See *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 980 F.3d 157, 163 (1st Cir. 2020).

²⁵⁶ See, e.g., *Brown v. Bd. of Educ.*, 347 U.S. 483, 495 (holding that it does not matter for Equal Protection purposes whether segregated schools are “equal” in quality).

race.²⁵⁷ It does not matter, on this view, whether affirmative action helps or harms historically disadvantaged groups. By implementing any race-based preferences at all, affirmative action commits the very moral sin it purports to cure.

Big Data Affirmative Action is not vulnerable to any of these critiques. The response to all of them is the same: Big Data Affirmative Action policies do not operate as unearned windfalls to members of minority groups. Instead, they are precise *correctives* for discrimination that those individuals would otherwise suffer. In this sense, Big Data Affirmative Action is best analogized to civil damages for discrimination—not, say, a lottery open only to members of historically disadvantaged groups.

If a brilliant Black woman is hired by a white-shoe law firm, and her firm has a Big Data Affirmative Action hiring policy, what should her colleagues think of her? It's simple: they should think that she was hired under *fair* conditions of *nondiscrimination*. Yes, the firm added some bonus points to her application file. But so what? Those are just the points she would otherwise have been *denied* because of her race. Surely, no one thinks that the firm *ought to* have discriminatorily withheld points. Thus, no one should think less of her if it gives them back.²⁵⁸

The same logic applies to the “more for you means less for me” critique. It is true that, in zero-sum settings, Big Data Affirmative Action might cause members of some groups to lose a benefit when members of another gain it. But this, again, is objectionable only if the losers would have gotten the benefit under *nondiscriminatory* decision procedures. No one is entitled to the advantage of a decision-maker's discrimination against their competitors. Since Big Data Affirmative Action policies operate to carefully *remove* discrimination from decisions, those who “lose” under such systems generally have no cause to complain.²⁵⁹

²⁵⁷ See ANSELL, *supra* note 254, at 136.

²⁵⁸ Skeptics might further argue that Big Data Affirmative Action does not give the attorney precisely the number of points she was denied because of her race. Rather, the adjustment represents a best statistical guess for someone like her, around which there remains some individual variance. True. But this means that she might just as well have gotten an adjustment that was too small as too large. Her colleagues have no reason to suspect one rather than the other and thus no reason to assume she is any more or less deserving of admission than they are.

²⁵⁹ For a more plausible version of this objection, consider a context in which two groups suffered discriminatory harm, but the Big Data Affirmative Action remedy applied to only one of them. Then the excluded group might have a kind of normative claim against the employer: Why should the included group's injury be remedied, but not theirs? It is an interesting and, to my knowledge, open question whether this normative claim would have legal weight. Certainly, the excluded group might have a legal claim for the first-order discrimination that they faced. But should they have a claim based on the differential *remedy*? Should we penalize employers for fixing some of their mistakes, simply because

Finally, there is the argument that assigning value based on race is simply wrong, regardless of whether the value is positive or negative. Even accepting the premise, Big Data Affirmative Action does the *exact opposite*. Big Data Affirmative Action policies are predicated on evidence that *someone else*—an employer, judge, doctor, etc.—has assigned (negative) value to race. The policy then *undoes* that assignment by adjusting the outcome to be race neutral. It is not morally objectionable if an intervention to correct racial discrimination ends up being race based. That is just how remedies work; they flow to the injured. Big Data Affirmative Action, then, is not reverse racism at all. It is, on the contrary, a reversal of racism.

For better or worse, the same cannot be said of every other style of affirmative action policy. Campus affirmative action, to which the reverse racism label is often applied, is not designed as a remedy for discrimination. Recall that the legal justification for affirmative action in admissions is that colleges have a compelling interest in attaining diversity of thought on campus.²⁶⁰ Under that framework, affirmative action *does* assign value on the basis of race—membership in a racial minority group has positive diversity value.

Conservative critics may nevertheless be wrong about campus affirmative action. It is arguably *not* always wrong to assign value based on race. Instead, whether it is wrong may depend on many factors, like the amount of value being assigned, the reason for assigning it, and the downstream effects of doing so. But this response is complex and perhaps debatable. Big Data Affirmative Action, working as it does to *undo* race-based value assignments, is justified irrespective of that debate.

B. *Individual Injuries, Average Remedies*

As just argued, Big Data Affirmative Action is not reverse racism because it supplies legitimate remedies for legitimate injuries. But how well do the remedies match the harms? Consider, for example, a given statistical model for measuring race-based salary discrimination that identifies a 10% pay gap between Black and white employees at a firm. A Big Data Affirmative Action policy based on that model would automatically adjust Black employee salaries up by the same amount. But—one could object—some Black employees surely suffered more discriminatory harm than that—and some less. Big Data Affirmative Action therefore supplies only

they did not fix all of them? This is a harder question. At any rate, such considerations show that, when institutions promulgate Big Data Affirmative Action policies, they should try to make those policies comprehensive.

²⁶⁰ See *supra* Section III.A.1.

“average”²⁶¹ remedies that will often either over- or undercorrect individual harm.

To begin, this factual characterization of Big Data Affirmative Action is not quite right. Neither statistical models measuring discrimination nor the Big Data Affirmative Action policies based on them need treat every Black employee the same. As discussed above, in the statistical models, race can be intersected with other features to measure intragroup variation in discrimination.²⁶² Thus, a Big Data Affirmative Action policy could supply different benefits to Black women than Black men, assuming they were treated differently by the relevant institution. Or it could vary its benefits by both race and age. Or by both race and education. And so on.

Adding this kind of intragroup variation to Big Data Affirmative Action policies blunts the force of the “average remedies” critique. But it does not refute the argument entirely. No amount of statistical refinement can produce a policy that *perfectly* corrects *every* discriminatory injury. There will always be some amount of averaging—say, across older Hispanic women who do engineering work.

This, however, is not a good reason to refuse to implement Big Data Affirmative Action. The objection here sounds in error rates: any Big Data Affirmative Action policy will compensate some people a bit too little for their harm and some a bit too much. But without Big Data Affirmative Action, the alternative would most often be no recovery at all. Absent affirmative action, litigation is the likely alternative path to compensation. And, as discussed above, the evidentiary burdens of Title VII and similar laws render many discriminatory injuries—indeed the very ones at which Big Data Affirmative Action is aimed—irremediable.²⁶³ Big Data Affirmative Action thus introduces small individual errors in order to avoid the massive and pervasive error of a uniform no-recovery rule.

This trade-off is not unique to Big Data Affirmative Action. On the contrary, for these same reasons, averaged remedies are commonplace everywhere in antidiscrimination law—including both individual and class litigation. Consider a class action alleging that some employer’s policy imposes a disparate racial or other impact. There, the employer is liable for *average*, not *individual*, impacts on protected groups.²⁶⁴ Remedies, too, are often averaged. An injunction against a challenged policy is a one-size-fits-all remedy. It benefits the majority of class members who were harmed by

²⁶¹ Where “average” means something like “least squares regression coefficient,” as opposed to “population mean.”

²⁶² See *supra* notes 58–59 and accompanying text.

²⁶³ See *supra* Section III.C.

²⁶⁴ See Salib, *supra* note 57, at 535–36.

the enjoined policy, but it harms those few who may have benefitted from it. Damages in class actions, too, are usually awarded on an averaged basis. In the presence of small but variable harms, it would be woefully inefficient to force every single class member to stage a mini-trial to prove his or her personal losses. Instead, class members are routinely divided into subpopulations and awarded damages based on features like job function, length of tenure, protected group, and the like. This practice closely mirrors the above-described strategies for inducing intragroup variation into Big Data Affirmative Action policies.

Even individual antidiscrimination suits involve substantial averaging. One way to establish individual Title VII liability is for a minority plaintiff to show that she possessed the requisite professional qualifications for a job, that she was rejected, and that the employer continued seeking applicants with the same professional qualifications.²⁶⁵ That is, she may show that, on average, a candidate like her ought to have been hired. The employer may rebut this showing with individualized evidence that the plaintiff was actually a less desirable candidate than an average candidate like her.²⁶⁶ But depending on the records available, such evidence may be difficult to produce, or the factfinder may reject it as pretext.²⁶⁷ That is, the average may carry the day in establishing liability. Averaging is even more important at the damages phase. If an employee has proved she was underpaid for discriminatory reasons, the only evidence of what she is owed is what others like her, on average, were paid. Because she was discriminated against, counterfactual information about what she personally would have been paid does not exist.

Thus, all antidiscrimination law authorizes average remedies for individual harms. To be sure, it does so in varying degrees. The magnitude of errors from averaging in individual suits may be smaller than in class actions. But this is often justified because the realistic alternative to class actions is not individual suits, but rather no remedy at all—even for valid claims. Averaging across a class thus reduces error at the relevant margin.

So too for affirmative action. The whole point of affirmative action is to authorize voluntary, race-conscious remedies for discriminatory harms that would otherwise go unredressed. Here again, averaging is crucial. If the law required affirmative action policies to suss out highly granular individual discriminatory harm—and then apply highly individual remedies—no institution would implement them. As discussed above, under such a rule,

²⁶⁵ See *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 802 (1973).

²⁶⁶ See *id.*

²⁶⁷ *Id.* at 804.

the evidence required to support an affirmative action policy would *also* support a lawsuit against the implementing entity. The Supreme Court has therefore repeatedly insisted that such evidence is not required. Instead, averaged evidence—and thus averaged affirmative action remedies—must be allowed.

History bears this out. The affirmative action remedies that the Court has historically approved have included minimal or no individualization. As discussed above, *Swann* upheld a plan that bused *all* Black students to white schools. *Weber* and *Paradise* upheld quotas that benefitted Black craftworkers and officers without differentiating between different magnitudes of individual discriminatory harm. *Johnson* authorized the hiring of women with lower test scores than men. It did not require a complicated process by which test scores would be individually adjusted to account for variations in women's discriminatory harm. These, too, are average remedies. And blunt ones. By comparison, then, with historically approved affirmative action plans, the averaging errors inherent in Big Data Affirmative Action are much smaller.

C. Data Skepticism

Statistical analysis is not physics. One cannot expect to get exactly the same result from the same experiment every time one runs it. Datasets are imperfect. They include elements of randomness. They can be analyzed via different techniques under different assumptions. Given all of this, would Big Data Affirmative Action policies stand on firm empirical footing?

As argued above, the *legal* answer to this question is yes; Big Data Affirmative Action's statistical models are sufficiently reliable to satisfy the strong basis in evidence standard. But this does not quite answer the normative question. *Should* Big Data Affirmative Action be considered empirically sound enough to be implemented as an antidiscrimination strategy?

One can imagine two kinds of concerns here. First, some might worry that Big Data Affirmative Action will sometimes go too far. Perhaps certain policies, predicated on anomalous or unreplicable evidence would bestow large and unfair benefits in zero-sum contexts. At a first cut, there is little reason to worry about this. Certainly, some statistical models underpinning Big Data Affirmative Action would overestimate discrimination somewhat. The resulting policies would be modestly overgenerous. But for every good-faith *overestimation* resulting from random noise, we should expect a counteracting *underestimation*. Most policies will be close to correct, and, on average, they should be *very* close.

The alternative is to decline to implement Big Data Affirmative Action at all and to thus allow pervasive discrimination to continue unchecked. Even if one objects to errors in racialized treatment, why prefer this option? To do so, one must think that small—and likely zero, on average—errors *overcorrecting* discrimination are worse than the huge, blanket error of *no* correction. There is little sense in such a view.

In fact, the second kind of concern goes one step further. Perhaps empirical errors would not be randomly distributed, but would rather be the result of motivated reasoning, or even fraud. Here, too, though, one can anticipate errors on both sides of the ideal point. Maybe some overzealous progressive institutions would rely on statistical models that were too generous in granting affirmative action benefits. But plenty of cautious or conservative-leaning institutions could effectively entrench even larger opposing errors by refusing to consider Big Data Affirmative Action at all.

Moreover, when it comes to data analysis, it is not so easy to cook the books. A recent study by Professors Jonathan Masur and Eric Posner examines such attempts by Trump Administration agencies.²⁶⁸ It finds that, despite the agencies' best efforts, they were largely unable to produce credible cost-benefit analyses to justify rolling back Obama-era policies.²⁶⁹ Relatedly, in academia, empirical social scientists have developed techniques like "pre-registration," "open data," and "open code" to guard against improper research techniques.²⁷⁰ These techniques could easily be adopted for Big Data Affirmative Action as a prophylactic against charges of motivated reasoning.

In the end, it is always a mistake to treat data scientists as infallible oracles. But so too would it be a mistake to dismiss their techniques. Data analysis is a powerful tool for learning about what is actually happening in the world. And for problems like discrimination, that makes it a powerful tool for improving the world via well-calibrated policy. Moreover, statistical social science can be done transparently, allowing skeptics to check practitioners' work for both error and malfeasance.

²⁶⁸ See Jonathan S. Masur & Eric A. Posner, *Chevronizing Around Cost-Benefit Analysis*, 70 DUKE L.J. 1109, 1114–36 (2021).

²⁶⁹ *Id.* at 1136 ("The preceding examples demonstrate that [cost-benefit analysis] is not as malleable as some of its critics have contended.").

²⁷⁰ Erin Standen, *Open Science, Pre-Registration, and Striving for Better Research Practices*, AM. PSYCH. ASS'N (Nov. 2019), <https://www.apa.org/science/about/psa/2019/11/better-research-practices> [<https://perma.cc/W4BU-RQ2J>].

D. Weak Medicine

Suppose Big Data Affirmative Action policies *could* be trusted to redress particular kinds of discrimination by particular institutions. Would that be enough? Consider again criminal sentencing. The gross disparity in federal criminal sentences between Black and white defendants is roughly 64%.²⁷¹ But a Big Data Affirmative Action sentencing policy would control for causes other than discrimination in the courthouse. As a result, such a policy would probably reduce the average Black defendant's sentence only by around 9%.²⁷² That leaves a disparity of more than fifty percentage points untouched. What good is Big Data Affirmative Action if it constitutes such weak medicine?

There are two possible responses here—one pessimistic and one optimistic. The pessimistic response is to say that whatever its shortcomings, Big Data Affirmative Action is the best we can do to combat persistent discriminatory inequality. Proposals to fix discrimination by excising bias from human decisions simply will not work.²⁷³ By contrast, affirmative action corrects unfair outcomes directly. And Big Data Affirmative Action's careful, modest adjustments are the biggest ones our law allows.

The second, optimistic response is to say that Big Data Affirmative Action need not be weak medicine at all. It is instead like real medication—perhaps a painkiller. One capsule may be weak. But patients with an especially bad headache can take two. Patients recovering from surgery may be prescribed an ultra-strong dose by their doctor. What matters is not the strength of the individual pill, but rather the availability of enough of them to get the job done.

It is true that a Big Data Affirmative Action policy for federal criminal sentencing would eliminate only a small fraction of the racial disparity there. But that is because much of the total disparity is caused by inequalities upstream of sentencing. Factors like prior arrests and education—inputs into sentencing—are themselves deeply unequal.

Insofar as these upstream inequalities are also the result of discrimination, *they too* can be the target of Big Data Affirmative Action. The strong basis in evidence rule forbids one institution's affirmative action policy from fixing other institutions' discrimination. But it places no limit on the *total number* of institutions with affirmative action policies.

²⁷¹ Rehavi & Starr, *supra* note 17, at 1321.

²⁷² *Id.* at 1337–38.

²⁷³ See *supra* Section I.B. And even if such interventions did work, producing completely fair-minded individual decision-makers, the per-intervention effect would be no bigger than that of Big Data Affirmative Action. After all, Big Data Affirmative Action simply *simulates* fair-minded human decision-making.

Thus, the more widely Big Data Affirmative Action policies spread, the better. If disparities in arrests drive the disparity in sentencing, then reducing the former will substantially reduce the latter. Upstream of that, perhaps disparities in access to education drive disparities in arrests. Big Data Affirmative Action policies for educational institutions can reduce inequality there, which will in turn reduce the disparity in arrests, which will reduce the disparity in sentencing. And so on.

Thus, Big Data Affirmative Action, as a policy design, is an ecosystem, not a monoculture. Every individual policy, implemented at every individual institution, affects outcomes downstream of that institution. Each downstream institution may, in turn, implement its own Big Data Affirmative Action plan to redress the bias they would otherwise feed into the system. Every iteration moves the needle only a little bit toward justice, but small improvements eventually cascade into large social effects.

This, arguably, is how affirmative action is *supposed* to be implemented under our current legal rules. The strong basis in evidence standard carves the world's discrimination up into chunks. And it makes each chunk remediable by one—and only one—institution: namely, the institution that caused it. This eliminates the kind of “double recovery” that would occur if, say, a school *and* an employer both had policies correcting for the school's discrimination.²⁷⁴ But on the other hand, this division is empowering, inviting each institution to undo both the discriminatory harm it has caused *and* the resulting downstream effects.

In the end, then, Big Data Affirmative Action is weak medicine only if the patient takes too little. But with enough doses, administered carefully at the right time and place, it can be quite strong medicine, indeed.

E. Evasion and Adaptive Decision-Making

Would individuals whose decisions would be subject to revision by Big Data Affirmative Action accept such adjustments happily? And if not, would individual resistance be a serious practical barrier to implementing such policies?

Consider the three possible reactions to Big Data Affirmative Action: enthusiasm, disgruntlement, and indifference. Indifference will likely be by far the most common. After all, in all kinds of settings—hiring, school admissions, commercial transactions—individuals' decisions are *already*

²⁷⁴ I do not mean to suggest that such double recovery would necessarily be normatively objectionable. Only that it would surely be *objected to* by affirmative action's many critics. The current legal rule thus represents a compromise view of what affirmative action should do. It treats discrimination as a wrong to be remedied, but insists that remedies must be targeted and incremental, rather than broad and sweeping.

constrained by mediating bureaucratic processes. What would it matter to a hiring manager if, after he has negotiated a new salary with an employee, human resources—via an algorithm—adjusts that salary? Would he even know? How different would this be from the many other constraints on compensation—from budgets to benefits—that upper management already imposes? Alternatively, what does a lending officer care if, contrary to her recommendation, her employer approves a loan pursuant to a Big Data Affirmative Action policy? How different is that from the many other reasons—a bad home inspection, a renegotiated sale price, unexpected financial disclosures—for overruling her?

Big Data Affirmative Action, in short, would often be implemented where there are *already* numerous procedures and policies to shape and constrain individual decisions. For most people affected, then, one more procedure is unlikely to even register.

Moving on, progressives should be enthusiastic about Big Data Affirmative Action. It is, after all, perhaps the only legally sound, data-driven intervention that can make a real difference in reducing discriminatory harms. What is there to complain about?

Maybe some progressives would worry about complacency. Perhaps they believe that, in the past, they scrupulously interrogated their own decisions for bias and successfully excised it. Perhaps then, they worry, Big Data Affirmative Action's automated de-biasing would cause them to lose that vigilance. Maybe, lacking vigilance, they would start introducing more bias into the system than they had before—and more than the Big Data Affirmative Action policy anticipated. Then, the policy's benefits would be too small, and net discriminatory harms would increase. This story seems unlikely. As discussed above, it is extraordinarily difficult to durably de-bias one's own decision-making, either via mental effort or by therapeutic practice.²⁷⁵ Thus, individuals who believe that, before Big Data Affirmative Action, their own mental efforts are preventing biased decisions are probably just wrong. To the extent that they are not wrong, Big Data Affirmative Action's measures of discrimination—and thus its interventions—could and should be updated over time. If the implementation of a Big Data Affirmative Action policy led to some complacency, then version 2.0 could simply offset it.

This brings us, finally, to the disgruntled. Almost certainly, some members of some professions would bristle at the thought of an algorithm correcting their biased decisions. Such individuals might be especially prevalent in professions where good judgment is thought to be part and

²⁷⁵ See *supra* Section I.B.

parcel of the job—judges, for example. Or doctors. Even in these professions, it should be noted, one ought to expect much enthusiasm for and indifference to Big Data Affirmative Action. Plenty of judges are progressives who would welcome affirmative steps to reduce inequality in, say, criminal sentencing. And like corporate employees, judges are used to having their decisions changed ex post via bureaucratic processes like appeal and clemency.

What, however, would the objectors do? If they knew the amount by which a Big Data Affirmative Action policy would adjust their decisions, they could counteract it. They might shift their initial decisions in an equal and opposite direction. This would be a very bad look. Suppose a judge knew how much a Big Data Affirmative Action sentencing policy would reduce sentences for Black defendants. To counteract those adjustments, upon implementation of the policy, he would need to start *increasing* all Black defendants' sentences by the same amount. Such a reaction would be easy to detect,²⁷⁶ raising the specter of public embarrassment and, perhaps, Equal Protection challenges. In the end, the strategy would most likely fail.

Perhaps the judge could be subtler. Maybe he would increase sentences for Black defendants by *less* than the Big Data Affirmative Action adjustment. Or maybe he would vary his increases somewhat randomly. These strategies, too, would be obvious if his average sentence for Black defendants jumped after the introduction of the policy.

In any case, here, too, Big Data Affirmative Action policies could adapt to counteract such strategic behavior. If actors in an institution adapted to policy 1.0 by imposing harsher discriminatory penalties, then version 2.0 could simply be updated to match the new behavior. Such updates could even be targeted to the offending individuals, relying on the empirical measures economists already use to study interjudge variance in sentencing.²⁷⁷ Certainly, disgruntled objectors could then adjust their own behavior even *more* to overcome the new policy. But this is an arms race that Big Data Affirmative Action is sure to win. At some point, the offending individuals' discriminatory behavior would become so outrageous that some combination of law and bad press would surely shut them down.

An even better solution would be to reduce disgruntlement at the outset. This might be accomplished by decoupling, as much as possible, initial decisions from Big Data Affirmative Action adjustments. It is one thing for a judge to see his preferred sentence immediately overridden. It is another if

²⁷⁶ Even a cursory look at the data would show a large discontinuity in sentencing around the time of the policy's implementation.

²⁷⁷ See generally, e.g., Crystal S. Yang, *Have Interjudge Sentencing Disparities Increased in an Advisory Guidelines Regime? Evidence from Booker*, 89 N.Y.U. L. REV. 1268 (2014).

the reduced sentence is recorded only later, by the Bureau of Prisons or some other actor. Such decoupling does two things. First, it raises information costs for bad actors. It is hard to consciously counteract a policy when you are not sure exactly what the policy is doing. Second, such acoustic separation would likely reduce the subjective feelings of offense that judges and other high-status actors would experience at being overruled.

In the end, strategic adaptation of decision-making is unlikely to threaten the effectiveness of Big Data Affirmative Action. Most people whose decisions would be subject to such policies would carry on without change. Some would enthusiastically welcome such policies' benefits. And the few objectors to Big Data Affirmative Action would have few avenues for successful resistance.

F. *Are Humans Obsolete?*

Big Data Affirmative Action policies are designed to adjust otherwise-defective human decisions. But why involve humans at all? If, as argued above, human decision-makers are deeply biased and immune to all cures, why give them any hand in making hiring decisions, setting salaries, approving mortgages, determining prison sentences, or anything else of importance?

Why not instead let the very Big Data models needed to *correct* human decisions simply *make* those decisions in the first instance? A standard regression model of a company's salary decisions could yield a number of coefficients: each year of postsecondary education is worth X dollars to the company, each point on a work skills aptitude test is worth Y, and so on. Using these coefficients, one could simply input a given candidate's information into the model to calculate their salary.

To ensure that the salaries had no discriminatory bias, the statistical model could, as all Big Data Affirmative Action models do, calculate the effect of race. But when each individual's salary was calculated, the algorithm would be told that *everyone* was a member of the most favored group.²⁷⁸ The result would mimic the Big Data Affirmative Action interventions advocated here. Minority group members would, in effect, receive a "bonus" in the precise amount necessary to offset the penalty they would otherwise face on account of race.

Why not do this? Perhaps the main reason that institutions do not abandon human decision-making entirely has to do with the so-called "error

²⁷⁸ See Yang & Dobbie, *supra* note 99, at 346–48.

term” in statistical models.²⁷⁹ In the real world, traditional regression models essentially never capture *all* of the variables affecting outcomes in human decisions. There will always remain unexplained differences between applicants who are identical along all of the metrics the model accounts for. In the salary example, human decision-makers may mysteriously pay two candidates with the same education, test scores, work history, hometown, gender, race, etc. somewhat differently. By contrast, a totally automated decision system based on traditional, interpretable statistical models would treat these two candidates alike. Any two employees who looked the same, according to the model’s finite inputs, would receive precisely the same salary.

The question, then, is whether the inexplicable variance that comes with human decisions is worth preserving. Perhaps it is simply random, the product of irrelevant factors like what the decision-maker ate for breakfast. In that case, good riddance. But suppose the error term is attributable to some X-factor of human judgment—a first impression or a gut instinct. If the variance captured in the error term is the result of some important information that humans can observe—but computers cannot—it might be worth preserving. Then, a human-centered decision procedure might be necessary.

On the other hand, human-observed X-factors might be worse than random. They might be discriminatory. Suppose a business wishes to base salaries or hiring decisions in part on an applicant’s friendliness. But suppose that the evidence suggests that judgments about such factors are strongly biased along racial lines. What then?

One option is to throw out personality factors entirely. But there surely *are* more and less friendly people—even if humans systematically misevaluate this because of race. To wit, the same empirical studies that showed *some* racial bias in personality evaluations would likely *also* show that such evaluations were not *solely* based on race. Instead, humans’ friendliness assessments might measure something real and identifiable person-to-person, but with a noticeable racial skew. And if some people really are more friendly than others, that is a legitimate factor on which some employers might wish to base employment decisions.

Another solution, then, is to use human-driven friendliness evaluations, but eliminate their discriminatory skew using Big Data Affirmative Action. Here, humans could take friendliness into account in making their decisions. But the statistical model on which the Big Data Affirmative Action plan was

²⁷⁹ See Adam Hayes, *Error Term*, INVESTOPEDIA (Oct. 19, 2021), <https://www.investopedia.com/terms/e/errorterm.asp> [<https://perma.cc/54PW-QR7D>].

based would *not* use humans' measures of friendliness as a control. Remember, approaches like regression analysis can estimate the effect of race on salary, *all other factors equal*—where “all other factors” means “all other modelled variables.” If the model asks, “What is the effect of race, holding friendliness constant?,” but friendliness assessments *themselves* depend significantly on race, the model will substantially *underestimate* the effect of race.²⁸⁰ Conversely, a model that estimated race but ignored differences in friendliness estimates would capture in the race estimation the proportion of friendliness attributable to race. Then, the Big Data Affirmative Action adjustment would correct the *entire* human racial bias, including the proportion tied up in humans' impressions of friendliness. The portion of friendliness assessments *not* attributable to race would be shunted to the model's “error term” and thus left undisturbed.

Leaving friendliness out of the statistical model here sacrifices little. Remember, the point of the analysis is to figure out the effect of race on human decisions. The point is *not* to figure out the effect of friendliness. Controlling for nonrace variables serves to disaggregate the influence of legitimate, *nonracial* factors that nevertheless correlate somewhat with race. So, it is important to regress over variables that *correlate* with race, but which represent some other legitimate reason—say, job function—for a disparity. But it seems highly unlikely that *actual* levels of friendliness or collegiality—as opposed to humans' race-inflected *perceptions* of them—correlate with race at all. Thus, leaving such factors out of the statistical model avoids underestimating the effect of race but raises less risk of overestimating it.

It is worth noting here that the two goals at issue above—preserving variance but eliminating discrimination—can also be achieved while eliminating humans. To accomplish this, one could implement the two-stage algorithmic decision processes described in Section II.A.3. First a radically nonlinear model—like a deep neural network—would be trained to precisely mimic human decisions. Such algorithms can reproduce the kind of variance associated with human judgment, preserving the valuable information encoded there.²⁸¹ Indeed, doing so is their whole point.²⁸² Then, a second Big Data Affirmative Action model would intervene to adjust the first model's decision just enough to eliminate discriminatory bias.

²⁸⁰ See Alin, *supra* note 98, at 370–71 (explaining that the more interdependent variables in a model, the smaller the measured contribution of any given variable, because the other variables “contain much of the same information”).

²⁸¹ See Salib, *supra* note 57, at 535–36.

²⁸² *Id.*

In the end, the choice between using humans augmented by Big Data Affirmative Action and using nonlinear algorithms augmented by it comes down to preference. Neither approach produces decisions that are appreciably different from the other in terms of *either* accuracy or discrimination. Institutions attracted to the efficiency of automation may choose an all-algorithmic approach. And those that value the human touch may keep it.

G. Less Biased Humans or Less Biased Decisions?

Big Data Affirmative Action operates under a kind of fatalism. The core idea is that humans, despite their best efforts, simply cannot or will not stop discriminating. Thus, the best we can do for now is to *let them* keep discriminating and enlist algorithms to fix their bad decisions on the back end. Is this attitude a betrayal of basic antidiscrimination principles? And does the law allow it?

*Connecticut v. Teal*²⁸³ suggests, at first blush, that this approach is disfavored. There, an employer used an exam as the first step in its promotion process; a passing score was the minimum qualification for further consideration.²⁸⁴ The test had a disparate impact. The passage rate was 80% for white employees and 54% for Black ones.²⁸⁵ The employer then made up for this disparity by treating Black employees *better* than white ones at the *later* steps in the promotion process—a putative “affirmative-action program.”²⁸⁶ As a result, a *greater* proportion of Black employees than white employees were ultimately promoted. Yet the Supreme Court held that this “bottom line” result was no defense, and the Black employees could state a Title VII claim.²⁸⁷

There is, however, a crucial difference between Big Data Affirmative Action and the approach that the *Teal* Court rejected. Namely, the employer in *Teal* sought to remedy the harm its discriminatory test did to certain Black employees by helping *different* Black employees. Recall that a passing score on the test was the first, bare-minimum requirement to advance in the promotion process. Yet the employer’s affirmative action policy applied only at the later stages. The policy’s favorable treatment therefore did no good for the Black employees who were weeded out at the testing phase.

Seen this way, the Court’s decision seems eminently sensible. As it wrote, “It is clear that Congress never intended to give an employer license

²⁸³ 457 U.S. 440 (1982).

²⁸⁴ *Id.* at 443.

²⁸⁵ *Id.* at 443 n.4.

²⁸⁶ *Id.* at 444.

²⁸⁷ *Id.* at 444, 452.

to discriminate against some employees on the basis of race or sex merely because he favorably treats *other* members of the employees' group."²⁸⁸ Of course the Black employees who suffered discriminatory harm and got no remedy had cognizable claims.

With this framing in mind, *Teal* becomes compatible with Big Data Affirmative Action. Such policies would not justify discrimination against some by bestowing special benefits on others. Rather, the benefits—ex post adjustments to human decisions—would accrue to *the very same people* who stood to suffer discrimination. Big Data Affirmative Action is not aimed only at the bottom-line goal of equality of representation, though it could achieve that if broadly deployed. Instead, Big Data Affirmative Action functions also at the level of individuals. By undoing the effect of discrimination for the individuals who suffered it, Big Data Affirmative Action conforms precisely to the Supreme Court's vision of antidiscrimination policy. It promotes not just equality of outcomes, but "equality of opportunity and the elimination of discriminatory barriers."²⁸⁹

This answers the legal question. But what about the normative one? Here, the most satisfying response is to say that *of course* it would be best to eliminate discriminatory harms by eliminating discriminatory attitudes. However, we as of yet lack the tools to directly and permanently change human hearts and minds. In the meantime, discrimination persists, and people suffer. Surely, then, it is better to relieve that suffering via Big Data Affirmative Action—while continuing to research the psychology of discrimination—than to do nothing at all.

CONCLUSION

Racial—and other—discrimination remains one of the most persistent and pernicious problems in American life. It has deep roots, reaching back well before the nation's founding. And despite enormous legal progress in recognizing the rights of racial minorities, women, religious minorities, and others, discriminatory inequality continues to abound. Solutions are hard to come by. Well-intentioned interventions designed to improve human thinking and reduce biased decisions, unfortunately, produce scant improvements. And old-fashioned affirmative action policies are routinely invalidated as illegal. Big Data Affirmative Action offers a solution. It works directly at the locus of unfair disparities, eliminating them and producing nondiscriminatory results. And because of its empirical precision, Big Data Affirmative Action can satisfy the legal requirements that have lately

²⁸⁸ *Id.* at 455.

²⁸⁹ *Id.* at 449.

bedeviled less-sophisticated policies. Deployed carefully and broadly, then, Big Data Affirmative Action policies would gain traction on the most intractable of problems. In doing so, they would provide relief for the millions of Americans who suffer from pervasive, unfair, and discriminatory decisions affecting nearly every aspect of everyday life.