How Discrimination and Bias Shape Outcomes

Kevin Lang and Ariella Kahn-Lang Spitzer

Summary

In this article, economists Kevin Lang and Ariella Kahn-Lang Spitzer take up the expansive issue of discrimination, examining specifically how discrimination and bias shape people's outcomes. The authors focus primarily on discrimination by race, while acknowledging that discrimination exists along many other dimensions as well, including gender, sexual orientation, religion, and ethnicity. They describe evidence of substantial racial disparities in the labor market, education, criminal justice, health, and housing, and they show that in each of these domains, such disparities at least partially reflect discrimination.

Lang and Kahn-Lang Spitzer note that the disparities we see are both causes and results of discrimination, and that they reinforce each other. For instance, harsher treatment from the criminal justice system makes it more difficult for black people to get good jobs, which makes it more likely they'll live in poor neighborhoods and that their children will attend inferior schools.

The authors argue that simply prohibiting discrimination isn't effective, partly because it's hard to prevent discrimination along dimensions that are correlated with race. Rather, they write, policies are more likely to be successful if they aim to eliminate the statistical association between race and many other social and economic characteristics and to decrease the social distance between people of different races.

Kevin Lang is a professor in the Department of Economics at Boston University. Ariella Kahn-Lang Spitzer is a researcher at Mathematica Policy Research.

Jonathan Guryan of Northwestern University reviewed and critiqued a draft of this article.

We all use information about the groups to which other people belong so we can determine how to treat them. We may use such information consciously or unconsciously, and it may be based on accurate statistical inference or on inaccurate beliefs. Consequently, our treatment of others can depend on the groups with which we associate them. Using statistical information based on race or other observable factors to make inferences about a person can alter the behavior of either party, or of both: for example, students may respond to their teachers' low expectations by not working hard. In some cases, discrimination is necessary and efficient. But in other cases, it can create and maintain a system of inequality.

In this article we focus primarily on discrimination based on race. However, people discriminate based on far more things than just race, including gender, sexual orientation, religion, ethnicity, and almost any other observable characteristic. Even race itself is a somewhat ambiguous distinction, and one that has evolved over time. Though today many people conflate race with skin color, historically these concepts were distinct. For example, in the first half of the 20th century, southern and eastern European (SEE) immigrants to the United States were considered both white and racially inferior. This is reflected in low rates of intermarriage. In 1910–20, 86 percent of married second-generation Italian women age 18 to 33 were married to first- or second-generation Italian immigrants, a rate of endogamy that's higher than that among Asians and Hispanics today. Over time, in a process that has been described as "becoming white," the social distance between SEE immigrants and whites decreased, intermarriage rates increased, and SEE immigrants gradually were no longer considered racially distinct from whites. Some see Asians and Hispanics currently undergoing the same process.[1] In other words, even our conceptions of how we define race, and who belongs to which race, are shaped by societal perceptions and social norms.

In the United States, we have substantial evidence of racial disparities across many domains, including, but not limited to, the labor market, education, criminal justice, health, and housing. The evidence suggests that in each domain, at least some of the disparity is due to discrimination.

Figure 1. Mutual Reinforcement: Disparities, Identity, and Discrimination.

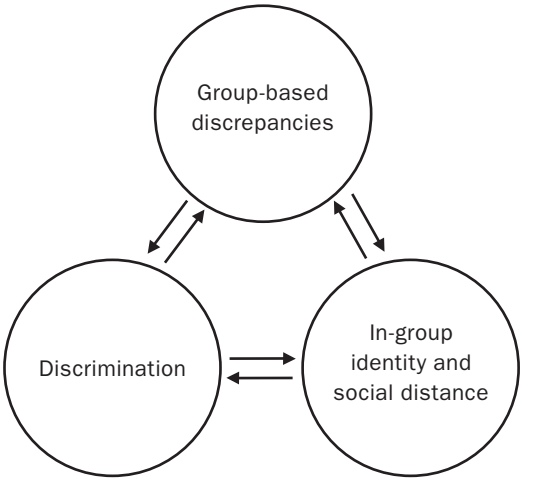

As figure 1 shows, disparities are both the cause and the result of discrimination. The fact that blacks are more likely than whites to come from disadvantaged backgrounds means that, on average, they arrive in kindergarten less prepared for school. This leads some teachers to have lower expectations for their

black students, which can produce further disparities in outcomes. Similarly, disparities both cause and are caused by racial identity. Racial disparities across domains contribute to residential and social distance between groups. Disparities between groups can be incorporated into a sense of within-race identity. Since group membership is important, our identities as a member of a group may also be important to us. We may generate a sense of wellbeing by confirming this identity. We may also be rewarded by other members of the group for confirming this identity, and punished or shunned if we don't. These group identities help create social distance between groups. Differences in language and norms of behavior can lead to miscommunication or to a reduced ability to assess members of a different group. This, too, can contribute to discrimination and reinforce disparities. And identity and discrimination reinforce each other. Identity contributes to the salience of race, and discrimination adds to social distance.

Disparities also reinforce each other. Harsher treatment by the justice system makes it harder for blacks to get good jobs or rent homes in better neighborhoods. This makes their children more likely to attend low-quality schools. And so on. This perspective suggests that ultimately eliminating discrimination would require reducing disparities across a large range of outcomes or finding key points of leverage.

We argue that policy designed to counter discrimination will be most effective when we consider the dynamics of discrimination. Simply prohibiting discrimination doesn't stop it, partly because we can still discriminate based on factors correlated with race. Furthermore, prohibiting discrimination based on factors correlated with race can actually increase race-based discrimination. Instead, we argue that the most effective policies are those that decrease disparities in outcomes and reduce residential or social distance between races. Increased spending on education for disadvantaged children and more integrated housing and schools are plausible candidates for this sort of focused policy.

What Is Discrimination?

We distinguish between *prejudice*, which refers to tastes, preferences, or inaccurate beliefs, and *discrimination*, which refers to actions or outcomes. Someone who dislikes working for a female supervisor is prejudiced. So is someone who has incorrect statistical beliefs, such as that immigrants are more likely than natives to commit felonies. But if these people don't change their behavior as a result of their prejudice, they aren't discriminating. To discriminate is to actively treat someone differently based on characteristics such as race, gender, ethnicity, or sexual orientation. Moreover, such behavior may not result in a *discriminatory outcome*. If the subjects of the discrimination can easily compensate for or avoid the discriminatory behavior, it is inconsequential. In some settings (but certainly not all), discriminatory behavior by a small number of people can be inconsequential and thus not lead to a discriminatory outcome.

Economists differentiate between *taste-based* (or *prejudice-based*) and *statistical* discrimination. Taste-based discrimination occurs when we treat people differently because we dislike or have false beliefs about people with a certain characteristic—for example, if an employer hires male candidates more often than female candidates because she prefers working with men. In

contrast, statistical discrimination describes differential treatment of individuals based on statistically valid inferences made using group membership—that is, using observable characteristics such as race to make statistically valid inferences that affect how we treat different people.

Statistical discrimination is universal. You're more likely to give up your seat on the bus for an elderly lady than a tall and broad young man. This is based on the perception that the elderly lady needs it more. But you don't know this. The young man may be recovering from surgery. You're just going with the odds. In other words, you're discriminating statistically. We all use signals of dress and, for better or worse, age, sex, and race to make statistical inferences about people and to act on those inferences. Statistical discrimination is often both legal and socially acceptable. Insurance companies pay for routine mammograms for older women but not for older men because the risk of breast cancer among men is low, although not zero. However, whether insurance companies should be able to charge men and women differently for health, disability, life, and auto insurance is controversial and varies among states.

As a society, we also accept taste-based discrimination in some settings. Some state courts have used privacy considerations to permit women-only health clubs as an exception to the civil rights law prohibiting sex discrimination in public accommodations. We often accept taste-based discrimination intended to counteract other forms of discrimination. Though controversial, certain types of affirmative action are legal and broadly considered socially acceptable.

Making the distinction between statistical and taste-based discrimination can be useful, because understanding which form of discrimination is at work can help us identify potentially effective policy solutions. But even when we can identify discrimination, it's often impossible to distinguish between the two forms.

Discrimination may also be conscious or reflect implicit bias. People may not be aware of the associations they have—associations that may even contradict their expressed beliefs. Implicit discrimination, documented by much evidence from social psychology, occurs when we treat members of groups differently based on these associations.[2] Thus an employer may unintentionally choose white candidates because she unconsciously associates whiteness with greater intelligence.

Implicit discrimination could reflect dislike or the accurate or inaccurate use of statistical association. Many people who identify as African American also demonstrate implicit bias against their own group, so it's unlikely that the bias is solely dislike. On the other hand, since implicit bias against blacks is more common among whites, it's unlikely that everyone unconsciously uses statistical association in an unbiased way. It's plausible that implicit bias reflects statistical associations unfiltered by rational updating.

Finally, discriminatory outcomes may not be the result of discrimination on the part of a given individual, but rather the result of institution-level policy or practice. For example, a policy that enforces harsher mandatory sentences on possession of crack cocaine relative to cocaine powder disproportionally affects blacks relative to whites, and therefore leads to discriminatory outcomes. In some cases, these policies may have been designed by a discriminating individual or a group of discriminating

individuals. In other cases, discriminatory outcomes may be the unintended effects of the policy or practice. In either case, discrimination at the institutional level can lead to widespread discriminatory outcomes. This is also referred to as *systemic discrimination.* In this article we primarily focus on individual- and group-level discrimination, although many of the factors that favor discrimination by individuals also influence institutional behavior.

Where Do We Discriminate?

Identifying and measuring discrimination is notoriously challenging. We can measure disparities among groups in various settings, but such disparities don't necessarily indicate discrimination. Disparities could reflect differences in preferences, innate differences between groups, and/or unequal treatment that occurred before contact with a given institution. For example, the absence of women in the National Football League could reflect biological differences between men and women, or the fact that young girls have less access to youth football and therefore don't develop the necessary skills and interest. It seems less likely to be driven by the discriminatory exclusion of a substantial number of qualified women.

Disparities in one domain can reinforce disparities in another.

So how do we identify discrimination? We can't conduct randomized controlled trials where we randomly assign some people to be black and others white. We can *pretend* that some people are black and others white, but as we'll see, that's not quite the same thing. Social scientists have used a number of techniques to try to identify discrimination in a wide range of settings, such as the labor market, medical care, education, criminal justice, and consumer markets, including credit and housing markets.

In this section we provide a brief overview of this research in various domains.[3] But these settings don't work in isolation. It's important to note their interconnected nature, which means that disparities in one domain can reinforce disparities in another.

The Labor Market

Substantial research provides evidence of differences in labor market outcomes among members of different races and ethnicities. An author of this article, Ariella Kahn-Lang Spitzer, has shown that in 2010, black men were 28 percent less likely to be employed relative to white men, and those who were employed earned 31 percent less annually.[4] While this doesn't prove that there's discrimination, it suggests that the labor market is an important domain in which to consider it.

A number of researchers have tried to isolate discrimination by looking at how much of the racial wage gap can be explained by observable characteristics such as education, test scores, location, and so on. Their results depend on the data source and the observable characteristics considered, but nearly all of these studies find that, after accounting for such observable factors, a smaller but still significant wage gap remains.[5] A typical finding is that the hourly earnings gap between blacks and whites falls by a little over half when we control for age, education, and a measure of cognitive skill.[6] But this approach has substantial weaknesses that

limit its usefulness to measure the extent of discrimination. First, we may not be able to control for some worker characteristics that are correlated with race, such as measures of school quality. These are often missing from data sets and are imperfect when included, but black people are likely to have attended lower-quality schools. Second, we risk overcontrolling for observable factors. For example, the difference in the occupational distribution of blacks and whites may explain a substantial portion of the earnings differential, but it may also reflect discrimination.

Despite these concerns, it's interesting to note that black and white women who had similar family incomes when they were growing up have similar wages and hours as adults, and thus similar personal incomes. In contrast, even when comparing men from families with similar incomes, black men have notably worse labor market outcomes than do white men.[7]

Although we can't assign race randomly, we can randomly assign résumés to individuals of different races, chosen to look as similar as possible and trained to act similarly, and see whether they have similar rates of interview and job offers.[8] Such studies, called "audit studies," fairly consistently find that employers discriminate among candidates based on race. For example, when sociologist Devah Pager assigned pairs of auditors to apply for jobs in Milwaukee, she found that white candidates were more than twice as likely to receive a follow-up call as black candidates.[9] But critics of audit studies point out that despite attempts to match on dress and appearance, such studies may pick up differences between applicants, and there are concerns that the auditors unconsciously bias the results.

To answer these criticisms, University of Chicago economists Marianne Bertrand and Sendhil Mullainathan randomly assigned black- and white-sounding names to résumés they sent to firms that were advertising job openings. By design, as in the audit studies, applicant quality was unrelated to the implied race of the candidate, but applicants with black-sounding names were still less likely to be called for an interview (6.4 percent versus 9.6 percent).[10] There has been some dispute about whether this experiment captures the effect of race or of names. The black female name with the highest callback rate got more callbacks than the white female name receiving the least, and this difference has been larger in other studies.[11]

More significantly, neither audit nor correspondence studies tell us whether this form of discrimination is important. Workers don't apply to jobs randomly. If black applicants have information about which firms discriminate, they may be able to avoid those firms with very little loss in terms of their labor market outcomes. Alternatively, if applying is sufficiently easy, a lower success rate may be easily offset by a larger number of applications.

Although these studies can provide compelling evidence of discrimination, they don't distinguish between statistical and taste-based discrimination. In an effort to identify whether discrimination is driven by prejudice (taste), economists Kerwin Kofi Charles and Jonathan Guryan compared wage disparities across states.[12] They argue that because only a minority of potential employees are black, if prejudice is the dominant factor behind discrimination then its impact should be driven by how prejudiced the relatively

less prejudiced individuals are. Consistent with this hypothesis, the researchers find that disparities are largest in states with the highest percentage of blacks among their population and, at the same time, where the relatively less prejudiced individuals (the 10th percentile in the state) are nevertheless more prejudiced than their counterparts in other states. This suggests that at least some of the gap between blacks and whites is driven by prejudice.

Education

There's little question that a large achievement gap between black and white children—and a somewhat smaller gap between non-Hispanic whites and Hispanics—emerges by kindergarten. There's some ambiguity as to whether the black-white gap grows or remains constant as children age, at least partially because of ambiguity in how to scale test scores.[13] Purdue economist Timothy N. Bond and Kevin Lang, an author of this article, have shown that when scores are scaled to predict educational outcomes, the black-white test score gap is fairly stable between kindergarten and seventh grade, remaining in the vicinity of somewhat less than a full year of predicted education.[14]

A wide range of factors have been proposed to explain test score gaps, including school quality, home inputs, early childhood education, differences in innate ability, and differences in education quality between and within schools. One strategy to measure the effects of family background is to analyze what share of the black-white test score differential can be explained by observable characteristics. At least some of the studies taking this approach find that controlling for socioeconomic and home environment differences between blacks and whites fully explains the test score gap.[15]

That's not to say there's nothing to be done. Teacher expectations can also have meaningful impacts on their students' academic performance.[16] Students perform worse when they're assigned to teachers with lower expectations of their ability, even after controlling for students' abilities. Both black and white teachers have been shown to have lower expectations for their black students, and, indeed, their black students do perform worse than their white students.

At the same time, findings that the black-white test score gap is predictable based on scores in kindergarten—and is also largely, if not entirely, explained by socioeconomic factors—suggest that successful policy may target disadvantaged students generally, rather than being race-specific.

The wide range of factors that contribute to education disparities highlights the interconnected nature of discrimination. Neighborhood segregation, which may in turn reflect housing discrimination or disparities in the labor market, can lead to segregated schools that may have fewer resources than predominantly white schools. Labor market discrimination can lead to socioeconomic status differences, which lead to fewer home resources for children. Another factor that contributes to the racial gaps in educational achievement is *stereotype threat*, the phenomenon by which individuals internalize stereotypes about the groups they belong to, and these beliefs become self-fulfilling. For example, a number of studies have found that when female students are reminded of their gender before a math test, their performance goes down. Similarly, one study found that black students performed worse on a test when it was described as an aptitude test, compared to when it was described simply as a problem-solving task.[17]

The fact that blacks, on average, receive less education than whites obscures a lesser-known fact: if we compare black and white students with similar high school test scores, the blacks go on to get more education than the whites. The disparity in educational attainment is driven by blacks' lower test scores, which, as we've noted, are found very early on. Lang and Boston University economist Michael Manove argue that blacks' greater investment in education, given their test scores, reflects their greater need to signal their ability. The authors further argue that this greater investment isn't driven by affirmative action in higher education: only very high-performing black students go on to colleges that use affirmative action, but the difference between blacks and whites is found primarily in the middle of the test score distribution. Thus, compared with a similar white student who would leave school after obtaining a high school diploma, a black student is more likely to attempt an associate degree. We can only speculate about whether this contributes to low completion rates. The twin result of lower test scores and higher educational attainment conditional on test scores suggests that improving cognitive outcomes for blacks through early interventions is likely to be a key way to reduce disparities.

Racial discrepancies are a particular issue in higher education. In 2017, black and Hispanic high school graduates were 16 and 3 percent less likely than white high school graduates, respectively, to attend college.[19] Furthermore, only 38 percent of black enrollees and 46 percent of Hispanic enrollees graduate from college within six years, compared to 62 percent of white enrollees.[20] This is partially driven by the fact that blacks, on average, attend colleges with lower graduation rates—which, in turn, at least partially reflects the preparedness of the students who attend them.

Elite colleges and universities have responded to racial disparities in education by favoring underrepresented groups in admissions decisions, commonly referred to as *affirmative action policies*. There's mixed evidence on how affirmative action affects college enrollment. By one estimate, these policies have nearly tripled the number of black students at elite institutions and more than doubled the number of Hispanic students.[21] But another study found that when California ended race-based admissions at the University of California, the impact on the student body's racial composition was minimal.[22] One explanation is that universities changed their admissions rules to consider factors that served as proxies for race.[23] Critics argue that affirmative action policies could be harmful if the students who are accepted due to the policy are less prepared than their classmates, and therefore struggle academically. Some evidence suggests that within-institution graduation rates for blacks rise when affirmative action is halted.[24] But multiple studies have found that affirmative action raises the overall graduation rates of blacks and Hispanics.[25]

Housing

Recent research demonstrates that neighborhoods are important, especially for young children.[26] Neighborhoods vary greatly with respect to safety, amenities, peer characteristics, public transportation, and access to job opportunities. And many resources, including public schools, are

distributed at the neighborhood level. Clearly, residential segregation contributes to the persistence of racial disparities. In the 2010 Census, the average white respondent lived in a neighborhood that was 75 percent white, 8 percent black, 11 percent Hispanic, and 5 percent Asian. In contrast, the average black respondent's neighborhood was 35 percent white, 45 percent black, 15 percent Hispanic, and 4 percent Asian.[27] Similarly, fewer than 5 percent of black children grow up in areas where the poverty rate is less than 10 percent and more than half of black fathers are present, while 63 percent of white children grow up in such low poverty areas with at least half of white fathers present. One study estimates that up to 25 percent of the gap in intergenerational mobility between blacks and whites is due to neighborhood effects, although those effects vary somewhat by race.[28]

Neighborhoods vary greatly with respect to safety, amenities, peer characteristics, public transportation, and access to job opportunities. And many resources, including public schools, are distributed at the neighborhood level.

Residential segregation is driven, at least in part, by discrimination. Audit studies—similar to those described above in the Labor Market section—reveal that prospective renters and buyers are treated differently depending on race, although the differences have generally declined over time.[29] Audit studies of mortgage applications would be illegal, but statistical comparisons suggest that blacks are less likely to receive a mortgage loan than whites with similar backgrounds.

Despite this evidence, the degree to which discrimination explains residential segregation is still uncertain. To some extent, residential segregation may reflect preferences on the part of most people to live near others of the same race. Furthermore, residential segregation can arise even if nobody prefers fully segregated neighborhoods. If whites are willing and able to pay more than blacks are to live in heavily white neighborhoods, we may end up with completely segregated neighborhoods even when everyone prefers some level of integration. If, for example, all whites preferred neighborhoods that are 20 percent black while all blacks preferred neighborhoods that are 40 percent black, we could still end up with every neighborhood being either all white or all black.[30]

The Justice System

Racial disparities exist at almost every level of the justice system. Black and Hispanic Americans are more likely to be arrested, less likely to be released on bail, and likely to receive harsher sentences than white Americans. Black adults are 5.9 times more likely to be incarcerated than are white adults; Hispanics are 3.1 times more likely.[31] Arrested black and Hispanic youth are also much more likely to be booked than are arrested white youth. Some of this variation can be explained by the characteristics of the offense and the suspect's prior record. But there's ample evidence that much of the discrepancy is due to differential treatment. For example, the American Civil

Liberties Union estimates that blacks are 3.7 times as likely as whites to be arrested for marijuana possession, despite the fact that the two groups have comparable rates of usage.[32] Likewise, having stopped a motorist, police are more likely to search the car of a driver whose race is different from their own. Since police officers are disproportionately white, this fact disadvantages blacks and Hispanics.[33]

Some of the differences in criminal justice outcomes can be explained by neighborhoods. Booking rates are higher in heavily black and Hispanic areas. Police presence, arrests, and bookings are more common in high-crime neighborhoods, where blacks and Hispanics are overrepresented.[34] Because of where they live, white juveniles are less likely to be caught when they commit a crime, less likely to be arrested if they are caught, less likely to be booked if they are arrested, and less likely to have a record if they're caught again.

Substantial evidence demonstrates direct discrimination and racial prejudice in criminal justice settings. One study examined felony trials in Florida's Sarasota and Lake counties, using random variation in the jury pool (not the actual jury) to examine the effect of race on convictions. All-white jury pools convicted 81 percent of black defendants but only 66 percent of white defendants, while pools with at least one black person were likely to convict whites and blacks equally.[35]

Some recent research is predicated on the argument that if criminal justice officials are prejudiced, if a black person and a white person are treated similarly (for example, they receive the same bail), then the black person should, on average, have a more positive outcome, such as being more likely to make a court date. One study finds that black defendants are less likely to be released on bail than white defendants who have the same estimated likelihood of reoffending, but that white defendants who are on the margin for release are 22 percent more likely to be rearrested prior to the outcome of the court case.[36] Similar studies document racial prejudice in traffic stops and death sentences.[37]

Not all studies point to discrimination in this direction. Some studies have found that when black and white motorists are stopped and searched, they're similarly likely to be caught with contraband, suggesting that the decision to search isn't prejudicial.[38] But the conclusion that the lack of difference in the conditional outcome implies no discrimination relies on strong assumptions about the distributions of probability of having contraband among blacks and whites. In a controversial study, Harvard's Roland Fryer Jr. finds that blacks and Hispanics are more likely to experience the use of force during a police stop but are no more likely than whites to be the victim of a police shooting if they're involved in an interaction. Moreover, when police officers do shoot, they're more likely to shoot first if the suspect is white than if the suspect is black.[39]

Disparities in Medical Treatment

Health is important, among other things, as an input into a child's educational success and later labor market success. A strong correlation exists between earnings and health, partially because higher earnings buy better medical care, but also because healthy individuals can earn more. There

are clear disparities in medical outcomes by race and ethnicity. In 2015 the Centers for Disease Control and Prevention estimated that the average life expectancy of a non-Hispanic black infant was 3.6 years less than that of a non-Hispanic white infant and 6.8 years less than that of a Hispanic infant. This finding is partially driven by the fact that infant mortality rates are more than twice as high for black babies than for white and Hispanic babies.[40] At the same time, unless we want to argue that the medical system discriminates in favor of Hispanics, the discrepancy between non-Hispanic whites and Hispanics tells us that we can't automatically ascribe the disparities to discrimination.

A substantial amount of research has documented differences in the medical care received by patients of different races. For example, black patients are less likely to receive such treatments as coronary artery bypass graft surgery, revascularization procedures, and thrombolytics. Some of these disparities reflect differences in where people live and the quality of their health insurance and, therefore, which physicians and hospitals they can access. It's hard to establish whether minorities and whites are treated differently when they have similar conditions and see similar physicians in similar hospitals. Economists Amitabh Chandra and Douglas Staiger argue that if there's discrimination in the provision of medical care, minority patients on the margin of receiving treatment should benefit more from treatment than their majority counterparts do. In contrast, they find that women and black patients realize slightly lower benefits from treatments following a heart attack, despite receiving less treatment. They argue that this suggests that, in fact, doctors may overtreat female and black patients due to equity and liability concerns.[41]

Why are black patients less likely to consume medical treatment? One answer is mistrust, driven by a long history of mistreatment of blacks by the medical field.

One explanation for inferior health outcomes for blacks has been lower usage of medical research and lower compliance with physician recommendations. Blacks are less likely to visit a doctor for either preventive care or treatment. In a study of patients being treated for chronic heart failure by the same physicians at the Veterans Health Administration, blacks were no less likely than whites to be prescribed the recommended medications, but were less likely than whites to comply with the physician's instructions. The study found that the strong adverse effects on blacks of failure to comply accounted for the racial disparity in survival probabilities.[42]

Why are black patients less likely to consume medical treatment? One answer is mistrust, driven by a long history of mistreatment of blacks by the medical field. Perhaps the most salient example is the Tuskegee Study of Untreated Syphilis, in which black men with syphilis were not informed of their diagnoses but were led to believe they were receiving treatment for a blood condition. Meanwhile, researchers passively observed the course of their untreated disease. Even after an effective syphilis treatment became available, the participants didn't receive access until

1972, when details of the study became public. After the details were released, black men in areas close to Tuskegee lowered their medical usage, causing a 1.5-year drop in their life expectancy.[43]

Better within-race communication may also explain blacks' lower take-up of medical care. One study set up a pop-up clinic providing preventive services in Oakland, CA, and randomized black men to a black or non-black (white or Asian) physician, about whom they were provided basic information, including a photo revealing the physician's race.[44] The patients were then offered complimentary cardiovascular screening and a flu shot. After viewing the photo, subject choices were independent of race prior to meeting with a physician, suggesting that the men weren't prejudiced against non-black doctors. But after an in-person meeting, subjects who met with a black physician were more likely to accept the services. Perhaps the black physicians were simply better doctors, or were more persuasive. Two results suggest otherwise. First, the subjects rated the black and non-black physicians equally highly on feedback forms. Second, the few subjects who didn't self-identify as black were less likely to choose the services when assigned to a black doctor. Moreover, black and non-black physicians spent similar amounts of time with the patients for the same services. The authors conclude that communication was better within-race, a finding that's reinforced by the beliefs of both blacks and whites that they communicate better with physicians of their own race.

The poorer communication between non-black doctors and their black patients and the lower rates of compliance by those patients may be causally related. There's often considerable uncertainty regarding the best treatment for a patient presenting a set of symptoms. If non-black physicians have more difficulty assessing the best form of treatment when working with a black patient, they're more likely to offer the treatment that they believe works best on average. This could result in overtreatment or undertreatment relative to whites, but in either case it will be worse, on average, than treatment that responds more precisely to the patient's condition. Since the treatment offered to black patients is, on average, less likely to be appropriate, the patient has less reason to comply. But, further, knowing that their black patients are less likely to comply with treatment, physicians may also shift their treatment recommendations to those that are less sensitive to imperfect compliance.[45]

Why Does Discrimination Persist?

In many ways, the persistence of discrimination remains a mystery to economists. If there are no true innate differences between groups, then there should be substantial returns to deviating and not discriminating. Thus, employers could make more profit by hiring more minority employees. We'd need only a small number of unprejudiced potential employers to eradicate discrimination. While dominant group members may benefit from maintaining their elite group status, it's implausible that group members could collude in a meaningful way. In this section we discuss some of the mechanisms that allow discrimination to exist and to persist over time.

Identity

One explanation for the persistence of discrimination is the role of group membership in our sense of identity.

According to this theory, individuals define themselves in the context of group membership. Social identification is defined as a "perception of oneness with a group of persons."[46] This identification leads people to identify with the characteristics, activities, and organizations associated with their group.

Crucially, identity includes a view of how people in the group should behave and a sense of who is not a group member. George Akerlof and Rachel Kranton suggest that individuals can, to some extent, reduce their otherness by adopting dominant-group behaviors, thereby increasing their social acceptability with the dominant group. But the reduction in otherness reduces the utility derived from behaving as dictated by "own-group" identity. Individuals choose how to behave depending on access to dominant group resources and own-group resources.[47]

In addition to any psychological benefits to identity, group membership offers important social and financial benefits. Groups provide companionship, as well as transfers from other group members. People are more altruistic towards own-group members and those they perceive as more similar to themselves. Economist Eli Berman discusses the many benefits of group membership in his study of ultra-Orthodox Jews in Israel.[48] Berman points out that group members benefit from access to social insurance, a substantial set of resources reserved for in-group members, community, and even potential marriage partners.

In the Akerlof-Kranton model, group membership is at least partially defined by conforming to the stereotypes of own-group identity. Berman describes this behavior as highly costly to ultra-Orthodox Jews in Israel, who may study full-time until age 40 while living in extreme poverty. Failing to conform to group norms, he explains, would result in loss of group membership. This may be more intuitive in the context of religion, but Fryer and Northwestern University's David Austen-Smith document similar mechanisms among black adolescents.[49] They argue that achieving academic success is viewed as "acting white" and therefore suggests disloyalty to a black identity. Using data on friendship networks and own-race friendships, Fryer and Austen-Smith find that a white student's popularity is positively correlated with GPA. In contrast, the correlation between GPA and popularity for black students is weak at low GPA levels and negative at high GPA levels. As the authors are well aware, these relationships are difficult to interpret. It's unlikely that getting good grades makes adolescent students more popular. More likely, it's the factors that are correlated with good grades that make students more popular. Still, either something weakens the link between these types of factors and popularity for blacks, or getting good grades ("acting white") actually reduces popularity for black students.

Not all research supports the acting-white hypothesis. Adjusted for socioeconomic background, blacks and whites report similar aspirations for completed education, spend similar amounts of time on homework, and have similar rates of aspirations except among the highest-performing students.[50] As we've already noted, the black-white test score gap at seventh grade is, if anything, less than would be predicted on the basis of kindergarten scores, suggesting that the factors causing poorer performance of blacks are largely socioeconomic rather than racial.

Although, as we've discussed previously with respect to race, the definitions of groups may change over time, the saliency of group membership can have important effects. First, we all use group membership to adjust our inferences about individuals. If police officers believe that blacks, on average, are more likely to commit crimes, they will treat otherwise similar blacks and whites differently; in turn, blacks will respond differently to police officers. The belief may even be self-fulfilling. Second, since group members often share common vernacular, cultural norms, and social circles, employers may find it difficult to identify high-ability out-group workers.[51] This encourages out-groups to shift toward easily observed forms of human capital investment and away from less easily observed forms.

Mutually Reinforcing Disparities

Disparities across domains, whether or not they're caused by discrimination, often reinforce each other. For example, if black children receive lower-quality schooling due to residential segregation and their parents' lower earnings, they will tend to be less prepared for the labor market. The correlation between lower skills and race can lead to statistical discrimination that perpetuates racial disparities above and beyond the intergenerational transmission of economic status that would occur in the absence of racial differences. Due to disparities that affect young children, such as socioeconomic status and residential segregation, black students, on average, arrive in kindergarten less prepared than their white counterparts. If this lowers teachers' expectations for their black students, the students may confirm these teachers' expectations. Lower school performance not only directly worsens labor market outcomes, but also contributes to statistical discrimination that further adversely affects employment and earnings. As we've said, these disparities can support a sense of racial identity and a view among whites that blacks are other, and thus they contribute to prejudice-based discrimination.

As long as there are salient racial disparities, people will use race as a heuristic to make statistical inferences about people. This creates substantial challenges to developing policies that effectively target discrimination.

Policy Implications

Simply prohibiting discrimination doesn't eliminate it, partly because it's hard to prevent discrimination along dimensions that are correlated with race. Furthermore, as long as there are salient racial disparities, people will use race as a heuristic to make statistical inferences about people, whether valid or invalid. This creates substantial challenges to developing policies that effectively target discrimination, and it suggests that eliminating the statistical association between race and many other social and economic characteristics must be both the goal of policy and the means by which that goal is achieved. As was the case with southern and eastern European immigrants in the early 20th century, doing so will likely come at the expense of elements of culture and identity. In this section we discuss

some of the challenges in designing policy to combat discrimination, and we highlight some promising policy directions.

Integration

While contact between people of different races could theoretically increase or decrease both taste and statistical discrimination, the clear weight of the evidence is that—at least as currently experienced—contact has desirable effects. A meta-analysis of 515 studies found strong overall support for intergroup contact theory, under which such contact tends to reduce prejudice.[52]

One study used the random assignment of freshmen to squadrons at the US Air Force Academy to examine how being assigned to a squadron with more blacks affects both attitudes and behaviors.[53] The study found that having an additional black member in a squadron of roughly 35 people increased the probability of having a black roommate as a sophomore (usually not a freshman squadron member) by about one percentage point, or about 18 percent. However, the authors' estimates imply that having one rather than no black squadron members has no effect on the probability of having a black sophomore roommate if the black member's academic score was substantially below the Air Force Academy average. Similarly, those exposed to more black peers with better admissions scores were more likely to say that they'd become more accepting of African Americans, and less likely to say they'd become less accepting. Again, it appears that attitudes were worsened only if the black squadron members to which they were exposed were substantially less academically prepared than average.

The Air Force Academy study suggests there's a role for policies that increase integration in schools and neighborhoods—at least if the policies don't involve the mixing of groups that are too disparate along other dimensions. The goal is to break down both prejudice and the statistical association between race and disadvantage. Increased exposure to blacks who confirm negative stereotypes is unlikely to be beneficial. At same time, the Air Force Academy experience suggests that fairly large gaps are still compatible with beneficial effects. Subject to legitimate concerns about projecting from results at a unique institution, these results imply that increasing the number of black students at an elite institution is likely to decrease prejudice, even if the average academic performance of the additional students is at the 25th percentile among the institution's enrolled students.

This perspective also sheds light on the debate about whether colleges should be required to use race-blind but not race-neutral policies instead of directly using race in admissions. As we've noted, by relying on correlates of race instead of race itself, colleges have largely been able to maintain the proportion of minorities that they enroll. However, such policies end up enrolling a quite different set of students, with noticeably lower test scores among black students.[54] In addition, when we use policies that are race-blind but not race-neutral, we may, intentionally or unintentionally, use selection criteria that favor certain groups. For example, a University of Texas rule admitting students who were in the top 10 percent of their high school class not only favors blacks and Hispanics but also favors students from schools with more low-performing students.

Information

As discussed, even when we decide it's not permissible to use characteristics such as race in decision-making, we sometimes permit policies that are race-blind but not race-neutral. This poses two key challenges. First, it requires identifying criteria that are necessary or reasonable for effective screening. Second, limiting information available to decision-makers sometimes has unintended consequences.

In *Griggs v Duke Power Co.*, a 1971 decision it has since overturned, the Supreme Court ruled that Duke Power Company's use of a high school diploma as a requirement for employment in certain jobs was illegal because it had an adverse disparate impact on black applicants, and because there was no business necessity, since many people without high school diplomas had done those jobs successfully. The Court's struggle with disparate impact cases reflects the difficulty of determining whether a practice is designed to discriminate or whether it exists for good reason, with disparate impact an unfortunate unintended consequence. In *Griggs*, the Court ruled that to be permissible, the policy generating disparate impact must be shown to be a business necessity. However, the Court later relaxed this standard.

The standard for justifying adverse impact has become increasingly relevant in the era of big data, when algorithms can predict an individual's race with great accuracy. But given that such algorithms are often quite opaque, an organization could unwittingly use race in its decisions. Consider a judge who, when deciding whether to grant bail, cares only about the probability that the defendant will be rearrested prior to disposition. The judge may turn to an algorithm that predicts the likelihood of re-arrest. Many predictors of re-arrest, such as number of prior offenses, zip code, and family situation, will be correlated with race. Therefore, any model will create a prediction that's correlated with race. Determining which variables should be included, and how much correlation is too much, is far from straightforward. And as we noted above, blacks who commit a crime are more likely than whites to be arrested. So even an ostensibly unbiased algorithm generated using arrest data will likely result in bias.

It can even be challenging to determine whether a policy limiting the information used in screening will decrease discrimination. Consider the heated debate around the use of criminal records to screen potential employees. We know that revealing a criminal record reduces the probability that a worker will be interviewed or hired. Blacks are more likely than whites to have criminal records. Therefore, requiring information about criminal records could be expected to have an adverse impact on blacks. In response, some cities and states have "banned the box," prohibiting employers from requesting such information on applications, though typically they can later conduct a background check or ask about it during an interview. But when firms can't use criminal records to screen applicants, they may screen by using correlates of criminal history, such as being a young black man. Rutgers University economist Amanda Agan and Sonja Starr, a law professor at the University of Michigan, submitted job applications for black and white male candidates with randomly assigned résumés both before and after the box was banned in New Jersey and New York City. Before "ban the box," companies with the box called black and white male

candidates without criminal records for interviews at similar rates. After the box was banned, callbacks of black men declined relative to otherwise similar whites.[55] Another study found that when states ban the box, the probability of employment falls by 3.4 percentage points for low-skill young black men and 2.3 percentage points for low-skill young Hispanic men.[56]

Other studies have consistently found that increasing information in hiring tends to help black candidates. Abigail Wozniak, of the Federal Reserve Bank of Minneapolis, finds that drug testing increases the employment of black candidates.[57] Like Agan and Starr, Wozniak hypothesizes that when employers have less information, they rely more on race to make decisions. Similarly, there's evidence that requirements for occupational licensing increase the share of minority workers in an occupation, despite their lower pass rates on licensing exams.[58] Another study finds that prohibiting employers from using credit reports in hiring reduced job-finding rates for blacks.[59] These studies show that adding more information in hiring, even when that information is highly correlated with race, may actually move employers away from using race directly. This suggests a more challenging job for the courts when judging adverse impact cases, further complicated by the fact that hiring practices by one firm tend to impact other firms. Therefore, if a set of firms introduces additional information into hiring practices, the quality of hired workers will likely increase, which may decrease the quality of the available worker pool. More research is needed to understand how these practices impact hiring when this feedback from the practices of individual firms to the pool of job seekers is taken into account.

Leverage Points

We've noted that disparities tend to reinforce each other. University of Washington sociologist Barbara Reskin describes these interrelated disparities as a "system of discrimination," arguing that policies aiming to counter discrimination in one domain must recognize this system or set of interactions. Consequently, she maintains that the most effective solution is a broad policy attack, like the civil rights reforms of the 1960s, that hits many components of the system.[60] Indeed, evidence suggests that civil rights policies reduced racial disparities in education and earnings.[61]

Reskin also argues for policies that target discrimination at *leverage points*—key points at which change is likely to have a substantial system-wide impact. She speculates that residential segregation is particularly promising as a leverage point because many resources exist at a neighborhood level, making disparities almost inevitable. A number of policies could help integrate neighborhoods. These include increasing the stock of affordable housing units in higher-income neighborhoods, reducing exclusionary land-use policies, and making public transportation more available and more affordable.[62]

Many researchers have pointed to education as a potential leverage point, although generally without using that term. Education, especially early education, is one of the first institutions we interact with in life. To the extent that education gives children the skills they need to be employable and successful as adults, any intervention that doesn't address the disparities in education is likely to be incomplete at best. We have strong evidence that high-quality early education has long-

term effects, and growing evidence that increasing school spending produces better student outcomes.[63]

Above we discussed the compelling research that teacher expectations affect student performance. Some evidence suggests that educating teachers about how their expectations affect student outcomes can reduce teacher bias.[64] There's also some evidence that black students who are randomly assigned to black teachers have higher rates of high school graduation and college attendance.[65] However, assigning black teachers primarily to black students might come at the cost of increasing segregation, especially if it's done by moving black teachers to classrooms with more black students rather than by increasing the supply of black teachers.

Conclusions

We've argued that the strong statistical relation between race and various outcomes fosters statistical discrimination, and that social distance reinforces this discrimination by interfering with within-race preferences and communication across races. Our two goals are therefore to reduce the statistical associations and the social distance. Unfortunately, as in the example of moving black teachers to primarily black neighborhoods, these goals can conflict. When possible, policies that can both decrease disparities and increase integration have the highest potential to decrease discrimination. And, as we've noted, a policy's impacts aren't always consistent with the intent of the policy. We therefore encourage policy makers to enable empirical research on potential policies whenever possible.

Endnotes

1. David Roediger, *Working towards Whiteness: How America's Immigrants Became White* (New York: Basic Books, 2005).

2. Marianne Bertrand, Dolly Chugh, and Sendhil Mullainathan, "Implicit Discrimination," *American Economic Revie*w 95, no. 2 (2005): 94–8, https://doi.org/10.1257/000282805774670365.

3. Readers looking for lengthier descriptions of empirical research on the existence of discrimination should refer to the following: labor market: Kerwin Kofi Charles and Jonathan Guryan, "Studying Discrimination: Fundamental Challenges and Recent Progress," *Annual Review of Economics* 3 (2011): 479–511, https://doi.org/10.1146/annurev.economics.102308.124448; education: George Farkas, "How Educational Inequality Develops," in *The Colors of Poverty: Why Racial and Ethnic Disparities Persist*, eds. Ann Chih Lin and David R. Harris (New York: Russell Sage, 2008): 105–34; criminal justice: Katherine J. Rosich, *Race, Ethnicity, and the Criminal Justice System* (Washington, DC: American Sociological Association, 2007), http://www.asanet.org/sites/default/files/savvy/images/press/docs/pdf/ASARaceCrime.pdf; housing and residential segregation: Leah Boustan, "Racial Residential Segregation in American Cities," in *Handbook of Urban Economics and Planning*, eds. Nancy Brooks, Kieran Donaghy, and Gerrit Knaap (New York: Oxford University Press, 2011); medical treatment: Institute of Medicine, *Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care* (Washington, DC: National Academies Press, 2003), https://doi.org/10.17226/12875.

4. Ariella Kahn-Lang, "Missing Black Men? The Impact of Under-Reporting on Estimates of Black Male Labor Market Outcomes," working paper, Harvard University, Cambridge, MA, 2019.

5. Patrick Bayer and Kerwin Kofi Charles, "Divergent Paths: A New Perspective on Earnings Differences between Black and White Men since 1940," *Quarterly Journal of Economics* 133 (2018): 1459–1501, https://doi.org/10.1093/qje/qjy003.

6. Kevin Lang and Michael Manove, "Education and Labor Market Discrimination," *American Economic Review* 101 (June 2011): 1467–96, https://doi.org/10.1257/aer.101.4.1467.

7. Raj Chetty et al., "Race and Economic Opportunity in the United States: An Intergenerational Perspective," working paper, National Bureau of Economic Research, Cambridge, MA, 2019, https://www.nber.org/papers/w24441.

8. Margery A. Turner, Michael Fix, and Raymond J. Struyk, *Opportunities Denied, Opportunities Diminished: Racial Discrimination in Hiring* (Washington, DC: Urban Institute Press, 1991).

9. Devah Pager, "The Mark of a Criminal Record," *American Journal of Sociology* 108 (2003): 937–75, https://doi.org/10.1086/374403.

10. Marianne Bertrand and Sendhil Mullainathan, "Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review* 94 (December 2004): 991–1013, https://doi.org/10.1257/0002828042002561.

11. Nicolas Jacquemet and Constantine Yannelis, "Indiscriminate Discrimination: A Correspondence Test for Ethnic Homophily in the Chicago Labor Market," *Labour Economics* 19 (December 2012): 824–32, https://doi.org/10.1016/j.labeco.2012.08.004.

12. Kerwin Kofi Charles and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *Journal of Political Economy* 116 (2008): 773–809, https://doi.org/10.1086/593073.

13. Timothy N. Bond and Kevin Lang, "The Evolution of the Black-White Test Score Gap in Grades K–3: The Fragility of Results," *Review of Economics and Statistics* 95 (2013): 1468–79, https://doi.org/10.1162/REST_a_00370.

14. Timothy N. Bond and Kevin Lang, "The Black-White Education Scaled Test-Score Gap in Grades K–7," *Journal of Human Resources* 53 (Fall 2018): 891–917, https://doi.org/10.3368/jhr.53.4.0916.8242R.

15. Bond and Lang, "Evolution."

16. Nicholas Papageorge, Seth Gershenson, and Kyung Min Kang, "Teacher Expectations Matter," working paper, National Bureau of Economic Research, Cambridge, MA, 2018, https://www.nber.org/papers/w25255.

17. Claude M. Steele and Joshua Aronson, "Stereotype Threat and the Intellectual Test Performance of African Americans," *Journal of Personality and Social Psychology* 69 (1995): 797–811, https://doi.org/10.1037//0022-3514.69.5.797.

18. Lang and Manove, "Education and Labor Market Discrimination."

19. National Center for Education Statistics, "Immediate College Enrollment Rate," *The Condition of Education* (2019), https://nces.ed.gov/programs/coe/indicator_cpa.asp.

20. Emily Tate, "Graduation Rates and Race," *Inside Higher Ed*, April 26, 2017, https://www.insidehighered.com/news/2017/04/26/college-completion-rates-vary-race-and-ethnicity-report-finds.

21. Thomas Espenshade and Chang Chung, "The Opportunity Cost of Admission Preferences at Elite Universities," *Social Science Quarterly* 86 (2005): 293–305, https://doi.org/10.1111/j.0038-4941.2005.00303.x.

22. Kate Antonovics and Ben Backes, "The Effect of Banning Affirmative Action on Human Capital Accumulation Prior to College Entry," *IZA Journal of Labor Economics* 3 (2014): article 5, https://doi.org/10.1186/2193-8997-3-5.

23. Kate Antonovics and Ben Backes, "The Effect of Banning Affirmative Action on College Admissions Policies and Student Quality," *Journal of Human Resources* 49 (2014): 295–322, https://doi.org/10.3368/jhr.49.2.295.

24. Peter Hinrichs, "Affirmative Action Bans and College Graduation Rates," *Economics of Education Review* 42 (2014): 43–52, https://doi.org/10.1016/j.econedurev.2014.06.005.

25. Sigal Alon and Marta Tienda, "Assessing the 'Mismatch' Hypothesis: Differences in College Graduation Rates by Institutional Selectivity," *Sociology of Education* 78 (2005): 294–315, https://doi.org/10.1177/003804070507800402; Kalena E. Cortes, "Do Bans on Affirmative Action Hurt Minority Students? Evidence from the Texas Top 10% Plan," *Economics of Education Review* 29 (2010): 1110–24, https://doi.org/10.1016/j.econedurev.2010.06.004; Hinrichs, "Affirmative Action Bans."

26. Raj Chetty, Nathaniel Hendren, and Lawrence F. Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review* 106 (2016): 855–902, https://doi.org/10.1257/aer.20150572.

27. John Logan and Brian J. Stults, *The Persistence of Segregation in the Metropolis: New Findings from the 2010 Census*, brief prepared for Project US2010, Brown University, Providence, RI, 2011.

28. Raj Chetty et al., "Race and Economic Opportunity in the United States: An Intergenerational Perspective," working paper, National Bureau of Economic Research, Cambridge, MA, 2019, https://www.nber.org/papers/w24441.

29. Margaret Turner et al., *Housing Discrimination against Racial and Ethnic Minorities 2012* (Washington, DC : Urban Institute, 2013).

30. Kevin Lang, *Poverty and Discrimination* (Princeton, NJ: Princeton University Press, 2007).

31. The Sentencing Project, "Regarding Racial Disparities in the United States Criminal Justice System," report of The Sentencing Project to the United Nations Special Rapporteur on Contemporary Forms of Racism, Racial Discrimination, Xenophobia, and Related Intolerance, 2018.

32. American Civil Liberties Union, *The War on Marijuana in Black and White* (New York: ACLU, 2013).

33. Shamena Anwar and Hanming Fang, "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence," *American Economic Review* 96 (March 2006): 127–51, https://doi.org/10.1257/000282806776157579.

34. Steven Raphael and Sandra V. Rozo, "Racial Disparities in the Acquisition of Juvenile Arrest Records," *Journal of Labor Economics* 37 (2019): S125–59, https://doi.org/10.1086/701068.

35. Shamena Anwar, Patrick Bayer, and Randi Hjalmarsson, "The Impact of Jury Race in Criminal Trials," *Quarterly Journal of Economics* 127 (2012): 1017–55, https://doi.org/10.1093/qje/qjs014.

36. David Arnold, Will Dobbie, and Crystal Yang, "Racial Bias in Bail Decisions," *Quarterly Journal of Economics* 133 (2018): 1885–1932, https://doi.org/10.1093/qje/qjy012.

37. Kate Antonovics and Brian G. Knight, "A New Look at Racial Profiling: Evidence from the Boston Police Department," *Review of Economics and Statistics* 91 (2009): 163–77, https://doi.org/10.1162/rest.91.1.163; Alberto Alesina and Eliana La Ferrara, "A Test of Racial Bias in Capital Sentencing," *American Economic Review* 104 (2014): 3397–3433, https://doi.org/10.1257/aer.104.11.3397.

38. John Knowles, Nicola Persico, and Petra Todd, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy* 109 (February 2001): 203–29.

39. Roland G. Fryer Jr., "An Empirical Analysis of Racial Differences in Police Use of Force," *Journal of Political Economy* 127 (2019): 1210–61, https://doi.org/10.1086/701423.

40. Elizabeth Arias and Jiaquan Xu, "United States Life Tables, 2015," *National Vital Statistics Reports* 67, no. 7, November 13, 2018, https://www.cdc.gov/nchs/data/nvsr/nvsr67/nvsr67_07-508.pdf.

41. Amitabh Chandra and Douglas O. Staiger, "Identifying Provider Prejudice in Healthcare," working paper, National Bureau of Economic Research, Cambridge, MA, 2010, https://www.nber.org/papers/w16382.

42. Emilia Simeonova, "Doctors, Patients and the Racial Mortality Gap," *Journal of Health Economics* 32 (2013): 895–908, https://doi.org/10.1016/j.jhealeco.2013.07.002.

43. Marcella Alsan and Marianne Wanamaker, "Tuskegee and the Health of Black Men," *Quarterly Journal of Economics* 133 (2017): 407–55, https://doi.org/10.1093/qje/qjx029.

44. Marcella Alsan, Owen Garrick, and Grant C. Graziani, "Does Diversity Matter for Health? Experimental Evidence from Oakland," working paper, National Bureau of Economic Research, Cambridge, MA, 2018, https://www.nber.org/papers/w24787.

45. Ana I. Balsa et al., "Clinical Uncertainty and Healthcare Disparities," *American Journal of Law and Medicine* 29 (2003): 203–19.

46. Blake E. Ashforth and Fred Mael, "Social Identity Theory and the Organization," *Academy of Management Review* 14 (1989): 20–39, https://doi.org/10.5465/amr.1989.4278999.

47. George A. Akerlof and Rachel E. Kranton, *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being* (Princeton, NJ: Princeton University Press, 2011).

48. Eli Berman, "Sect, Subsidy, and Sacrifice: An Economist's View of Ultra-Orthodox Jews," *Quarterly Journal of Economics* 115 (2000): 905–53, https://doi.org/10.1162/003355300554944.

49. David Austen-Smith and Roland G. Fryer, Jr., "An Economic Analysis of 'Acting White,'" *Quarterly Journal of Economics* 120 (2005): 551–83, https://doi.org/10.1093/qje/120.2.551.

50. Philip J. Cook and Jens Ludwig, "The Burden of 'Acting White': Do Black Adolescents Disparage Academic Achievement?," in *The Black-White Test Score Gap*, eds. Christopher Jencks and Meredith Phillips (Washington, DC: Brookings Institution Press, 1998), 375–400.

51. Kevin Lang, "A Language Theory of Discrimination," *Quarterly Journal of Economics* 101 (1986): 363–82, https://doi.org/10.2307/1891120.

52. Thomas F. Pettigrew and Linda R. Tropp, "A Meta-Analytic Test of Intergroup Contact Theory," *Journal of Personality and Social Psychology* 90 (November 2006): 751–83, https://doi.org/10.1037/0022-3514.90.5.751.

53. Scott E. Carrell, Mark Hoekstra, and James E. West, "The Impact of Intergroup Contact on Racial Attitudes and Revealed Preferences" working paper, National Bureau of Economic Research, Cambridge, MA, 2015, https://www.nber.org/papers/w20940.

54. Roland G. Fryer Jr., Glenn C. Loury, and Tolga Yuret, "An Economic Analysis of Color-Blind Affirmative Action," *Journal of Law, Economics, and Organization* 24 (2008): 319–55, https://doi.org/10.1093/jleo/ewm053.

55. Amanda Agan and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *Quarterly Journal of Economics* 133 (2018): 191–235, https://doi.org/10.1093/qje/qjx028.

56. Jennifer L. Doleac and Benjamin Hansen, "Does 'Ban the Box' Help or Hurt Low-Skilled Workers? Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden," working paper, National Bureau of Economic Research, Cambridge, MA, 2016, https://www.nber.org/papers/w22469.

57. Abigail Wozniak, "Discrimination and the Effects of Drug Testing on Black Employment," *Review of Economics and Statistics* 97 (2015): 548–66, https://doi.org/10.1162/REST_a_00482.

58. Marc T. Law and Mindy S. Marks, "Effects of Occupational Licensing Laws on Minorities: Evidence from the Progressive Era," *Journal of Law and Economics* 52 (2009): 351–66, https://doi.org/10.1086/596714.

59. Alex Bartik and Scott Nelson, "Credit Reports as Resumes: The Incidence of Pre-Employment Credit Screening," unpublished manuscript, University of Illinois at Urbana-Champaign, 2018.

60. Barbara Reskin, "The Race Discrimination System," *Annual Review of Sociology* 38 (2012): 17–35, https://doi.org/10.1146/annurev-soc-071811-145508.

61. John Donohue and James Heckman, "Continuous versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature* 29 (1991): 1603–43, https://www.jstor.org/stable/2727792.

62. Rebecca Diamond and Tim McQuade, "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy* 127 (June 2019): 1063–1117, https://doi.org/10.1086/701354; Thomas W. Sanchez and Marc Brennan, *The Right to Transportation: Moving to Equity* (Chicago: American Planning Association, 2007).

63. Jens Ludwig and Deborah A. Phillips, "Long-Term Effects of Head Start on Low-Income Children," *Annals of the New York Academy of Sciences* 1136 (2008): 257–68, https://doi.org/10.1196/annals.1425.005; C. Kirabo Jackson, "Does School Spending Matter? The New Literature on an Old Question," working paper, National Bureau of Economic Research, Cambridge, MA, 2018, https://www.nber.org/papers/w25368.

64. Papageorge, Gershenson, and Kang, "Teacher Expectations Matter."

65. Seth Gershenson et al., "The Long-Run Impacts of Same-Race Teachers," working paper, National Bureau of Economic Research, Cambridge, MA, 2018, https://www.nber.org/papers/w25254.