Implicit Bias as a Cognitive Manifestation of Systemic Racism

Manuel J. Galvan & B. Keith Payne

Explicitly prejudiced attitudes against Black Americans have declined gradually since the 1960s. Yet racial disparities and racial discrimination remain significant problems in the United States. How could discrimination and disparate outcomes remain constant even while racial prejudice decreased? Two prominent explanations have emerged to explain these puzzling trends. Sociologists have proposed that disparities and discrimination are perpetuated by systemic racism, or the policies, practices, and societal structures that disadvantage some racial groups compared with others. Simultaneously, psychologists have proposed that implicit biases may sustain discrimination even in the absence of explicit prejudice. In this essay, we explore newly discovered connections between systemic racism and implicit bias, how they challenge traditional views to reorient our understanding of implicit bias, and how they shed new light on strategies to reduce bias.

In 2022, artificial intelligence researchers at OpenAI released their latest development, ChatGPT. Using machine learning algorithms trained on large bodies of text, the chatbot could generate impressively human-sounding text responses on seemingly endless topics. Users soon began debating whether the technology had reached human-like levels of intelligence, even going so far as to invoke the concept of sentience.[1] Meanwhile, those with experience using artificial intelligence worried about a problem that has plagued the field for years: chatbots trained on human inputs are prone to saying racist, sexist, and otherwise offensive things.

The designers of ChatGPT had anticipated this problem with bias and had installed new filters to prevent the bot from saying inappropriate things. If you asked ChatGPT to tell a racist joke, for example, the bot would refuse, explaining: "I am not capable of generating offensive or harmful content." But as cognitive scientist Steven Piantadosi noted, those biased inferences were still there, and could be revealed by probing indirectly.[2] When he asked the chatbot to write a computer code function to check if someone would be a good scientist based on their race and gender, it generated code indicating that only white males were good scientists. When asked to create code to decide if a child's life should be saved based on their race and gender, the code indicated that the lives of Black males should not be saved.

https://doi.org/10.1162/daed_a_02051

As psychologists, we find that the continuing struggle to create artificial intelligence that is free from racism says more about humanity than about technology. The algorithms behind this chatbot make statistical predictions about what words go together, based on training with massive bodies of real-world text. When a statistical model returns a biased response, it reflects the biases in the human texts on which it was trained. Programmers can add rules like "don't say racist things," but that does not change the biases that are deeply embedded in the training environment. As a result, the chatbot may seem unbiased when asked directly but will reveal biases indirectly in countless ways. Artificial intelligence has thus encapsulated what psychologists have known about humans for decades: when a cognitive system that detects statistical regularities is immersed in an environment that is systemically biased, it will reproduce those biases.

The chatbot highlights something else about human psychology. When a robot reproduces biases, it is easy for humans to identify its environment as the source of the bias. Few people believe that there is something deep and essential about the robot's character that makes it racist. When humans form the same kinds of biased associations, however, people tend to attribute it to the attitudes, beliefs, or character of the person.

We argue that the human mind, like artificial intelligence, tracks statistical regularities in the social environment. When the mind is immersed in an environment of systemic racism, it tends to form biased associations and inferences about marginalized social groups. In fact, implicit bias is best understood as the cognitive reflection of systemic racism. This formulation may seem surprising: implicit bias has long been thought of as an individual trait or attitude, whereas systemic racism concerns structures, history, and social environments, rather than individuals. In this essay, we explore the connection between systemic racism and implicit bias: how it challenges traditional views to reorient our understanding of implicit bias, and how it sheds new light on strategies to reduce bias.

The theory of implicit bias grew out of efforts to understand gradual trends toward more egalitarian attitudes in standardized surveys. For example, beginning in the 1960s, white Americans have slowly caught up with Black Americans on issues of interpersonal discrimination. Today, over 90 percent of white and Black Americans support racially integrated schooling and reject laws against interracial marriage.[3] Another poll in 2019 found that 72 percent of white respondents believe it is never acceptable for a white person to use the N-word.[4] Such polling data illustrate the eventual decline in white people expressing explicit biases against Black Americans in surveys.

At the same time, actual racial disparities have remained largely undiminished. Relative to white Americans, Black Americans are far more likely to struggle with poverty, food insecurity, and unemployment.[5] Black Americans have 10 percent of the median net worth and half the median annual income of their white counterparts.[6] Such disparities are hard to address when racial discrimination persists. For decades, researchers have conducted field experiments responding to job postings with two versions of otherwise identical résumés: one with a name that implies a Black identity and the other with a name that implies a white identity. The rate of callbacks to the applicants is a measure of racial discrimination between otherwise equally qualified candidates. Recent meta-analyses of similar field experiments have demonstrated that racial discrimination in hiring has remained relatively constant since the late 1980s, and housing discrimination has decreased but remains potent.[7]

These trends created a puzzle. How could discrimination and disparate outcomes remain constant even while racial prejudice decreased? This question spurred innovations in thinking across the social sciences.

Sociologists developed the concept of systemic racism to account for the ways that inequalities can be perpetuated independent of individuals' attitudes and intentions. *Systemic racism* refers to policies, practices, and societal structures that disadvantage some racial groups compared with others.[8] This is distinct from more colloquial uses of the word racism to describe prejudicial thoughts, beliefs, or behaviors, which is often referred to as *interpersonal racism*. An essential theoretical contribution of systemic racism research is the recognition that individual actors do not simply act as racists or nonracists. For example, even if all discriminatory behavior stopped today, preexisting disparities in income, wealth, and educational opportunity would still ensure that racial inequalities are passed on to future generations.

Psychologists grappled with the puzzle of persistent discrimination and disparities amid attitudinal shifts toward egalitarianism by developing the concept of implicit bias. Implicit bias refers to positive or negative mental associations cued spontaneously by social groups. It is measured using cognitive tasks that test how those associations facilitate or interfere with task performance. Unlike survey methods, implicit tests are difficult to manipulate based on social desirability or norms against expressing prejudice. Studies suggest that implicit bias is widespread, even among people who explicitly endorse egalitarian attitudes.[9] If implicit bias leads to discriminatory behavior, it could explain the puzzle of widespread discrimination despite declining prejudice on surveys.

Implicit bias has traditionally been considered an individual attitude. Implicit bias tests and sequential priming tasks were developed as individual difference measures.[10] Most theories of implicit bias posit that implicit attitudes were learned from cultural biases early in development and became rigid because of immense repetition.[11]

The ideas of systemic racism and implicit bias were thus developed as very different approaches to solving the same puzzle. One focused on the ways that laws, policies, and social environments perpetuate inequalities without regard to individual attitudes. The other focused almost entirely on individual attitudes. However, recent research has reconsidered implicit bias as an individual trait. We argue that these two theoretical frameworks are not as different as was once assumed, and that implicit bias is in fact a cognitive reflection of systemic racism in the environment.

As research accumulated over the past two decades, several findings cast doubt on the individual-attitude view of implicit bias. For one, we expect individual attitudes to be stable over time. Implicit biases can be reliably detected in group averages using tests such as the Implicit Association Test (IAT) or the Affect Misattribution Procedure (AMP).[12] If one hundred randomly selected Americans completed the IAT, there is a very high likelihood that there would be a detectable average level of implicit bias across the group. However, longitudinal studies have found that while group averages are consistent over time, individual scores are quite unstable.[13] In other words, when a classroom of students takes the IAT, the rank order of the students will change such that the most and least biased students may not be the same when they retake the IAT at the end of the semester. Yet the classroom average will remain remarkably similar.

A second related puzzle is that average implicit bias does not change over the lifespan. Groups of younger and older Americans of various ages have been found to have very similar average implicit biases.[14] Under the traditional attitude assumption, this stability would naturally result from stable individual biases. However, given the temporal instability of individual-level bias, this age invariance is surprising. How can a variable that is unstable over two weeks be stable across a lifetime?

A third puzzling finding is that implicit biases of individuals are not strong predictors of individual discriminatory behavior ($r = 0.14$ to 0.24).[15] Yet when implicit biases are aggregated over larger geographic areas, they have much stronger associations to behavioral outcomes such as achievement gaps, disparities in shootings, health disparities, and internet searches using racial slurs.[16]

In light of these anomalies, psychologists B. Keith Payne, Heidi A. Vuletich, and Kristjen B. Lundberg developed the "bias of crowds" model to make sense of the large body of implicit bias research.[17] The basic assumption of the model is that implicit bias scores reflect the accessibility of concepts linked to social group categories. Concept associations can vary both chronically, as an individual difference from one person to another, and situationally, from one context to the next. For very stable individual constructs, like explicit racial attitudes, there is a lot of stable individual variation but little temporal variation within persons. Some people have explicit biases, others don't, but each individual's explicit biases are generally consistent over time. When stable traits are aggregated, the aggregate measure's stability is simply a reflection of the stability in individual differences in scores.

Despite its capriciousness at the individual level, implicit bias can be remarkably stable at the context level (such as city, county, or state level). We describe this as *emergent stability* because the aggregate stability cannot be reduced to stability of the individual scores. The ranking of people from highest to lowest implicit biases will shuffle over time, yet there will be a consistent mean level of implicit bias for the group. The bias of crowds model suggests that this consistent emergent stability reflects the relatively stable social context. Features of context make implicit associations between racial groups more or less prominent. When aggregation occurs, random variation at the individual level is reduced, enabling a clearer estimation of the influence of shared contextual factors on implicit bias. Given its emergent stability, implicit bias at the context level becomes a more theoretically and practically useful predictor and outcome for social scientists.

Many variables can be either measured as individual differences or averaged across individuals to measure contexts. For example, very stable aspects of personality have been found to vary across geographic regions. People in Middle America and the South (typically "red states") are more inclined to be "friendly and conventional," meaning they are higher in conscientiousness, agreeableness, and extraversion, and low in neuroticism and openness.[18] Regularities in regional personality structure are thought to be due to regularities in the physical environment, historical events, and cultural norms of the region.[19]

But unlike very stable personality constructs, implicit bias is limited as an individual difference variable, and is instead particularly powerful as a context-based measure. One reason, reviewed above, is that implicit bias scores are very low in stability. A second related reason is that implicit bias scores are highly context sensitive.[20] For example, experiments have shown that seeing Black Americans in a positively valenced context, like at church or a family barbecue, results in participants having lower anti-Black implicit biases compared with when they see Black Americans in a negative context, like prison.[21]

Because implicit bias is unstable and highly context sensitive, the average implicit bias in a city, county, or state is not reducible to the attitudes of the individuals that make up the context. This means that when we take a sample of participants from a given context to measure their implicit bias, the specific individuals in our sample are largely interchangeable. If you replaced the individuals sampled with another set of individuals from the same context, their aggregated scores would show the same level of implicit bias. Whatever influence is exerted by the context will be reflected in the scores of whoever inhabits those spaces. Because of this, aggregated implicit bias scores have proved to be extraordinarily sensitive indicators of systemic racism.

Much research suggests that implicit biases are influenced by contextual information in the environment. There is a large body of literature showing that implicit associations are influenced by experimental procedures. For example, one study attempted to influence the association between Middle Easterners and negative words by exposing participants to a slideshow showing Middle Eastern faces paired with positive images and white faces paired with neutral images. Relative to a control group shown the same images but without pairing the stimuli, the experimental participants showed a lessened degree of Middle Eastern implicit bias.[22] A meta-analysis of more than two hundred studies performed over many decades showed such evaluative conditioning effects on implicit biases are replicable, if small.[23] Other studies have demonstrated that counter-stereotypical experiences, such as positive interactions with a Black experimenter or reading about positive exemplars, can reduce negative implicit biases.[24] These laboratory studies demonstrate that implicit biases are subject to significant shifts due to environmental conditions. Conditions that match cultural stereotypes of marginalized groups strengthen associations with negative concepts and reinforce implicit biases.

Outside of the laboratory, systemic racism, as a set of long-standing structural, institutional, and cultural tendencies, has created the specific environmental conditions that would theoretically reinforce implicit biases. Most Black Americans are descendants of enslaved African people brought to the continent prior to the abolition of slavery.[25] The slave trade was a four-centuries-long brutal and dehumanizing regime that included capture, enslavement, destruction of African identity, disruption of families, and indoctrination of Black inferiority. Such trauma was also perpetuated by intergenerational familial trauma.[26] The legacy of slavery can be seen in contemporary patterns of distrust between ethnic groups, voting behavior, and cultural norms, belief, and values.[27] It also set the stage for the enormous wealth gap between white and Black people in the United States that has not meaningfully closed.[28]

While the Thirteenth Amendment officiated the end of slavery by federal law, there is a complex and sordid history between abolition in 1865 and the civil rights era in the late 1960s. In that time, systemic racism was a brazen and institutionalized set of practices that included Black Codes, sharecropping, lynching, Jim Crow laws, sundown towns, and redlining. Many studies have tried to estimate how these structures and events have shaped the contemporary context of Black-white inequality.

An analysis of U.S.-based health outcomes found that Jim Crow laws had an enduring impact on Black-white mortality rates from 1960 to 2009.[29] Southern counties with higher rates of historical lynchings from 1882 to 1930 had lower Black voter registration in modern elections.[30] Spatial proximity to sundown towns (that is, towns where Black people were subject to violence if they were present after sundown) predicts Black-white poverty disparities.[31] A number of studies have connected redlining, a legal practice until the passing of the Fair Housing Act in 1968, to current inequality. To provide only a sampling of recent research, historical redlining patterns are associated with life expectancy, the proportion of health care professionals, access to quality food, home heat vulnerability, environmental racism, cardiometabolic risk, tobacco retailer density, gentrification, alcohol outlet density, nonfatal shooting incidence, air pollution, fatal encounters with police, and COVID-19 exposure.[32]

These studies present strong empirical evidence that systemic racism has shaped the life outcomes of both Black and white populations in the United States. In other words, systemic racism is an important contextual factor that strongly influences who is successful, who has educational opportunities, who is exposed to violence and addiction, who lives in expensive homes and communities, and who languishes in poverty and within the carceral system. Such statistical regularities in our society are readily perceived as we walk to work, watch the news, or drive through segregated neighborhoods. For those of us in racially unequal regions of the country, which have been most impacted by systemic racism, there are myriad constant cues that one group has what the other group does not. The bias of crowds model suggests that the context of persistent and systematized inequality between racialized groups underlies the implicit associations we make between racialized groups and concepts like "good," "bad," "criminal," "smart," and "dumb" as measured by instruments like the IAT.

If implicit bias is an indicator of systemic racism, we would expect to find reliable associations between contextual aspects of systemic racism and implicit bias. Some studies consider which aspects of historical and current context might predict higher implicit bias in different geographic regions. As discussed previously, slavery has profoundly influenced current-day culture, behavior, wealth distribution, and other aspects of systemic racism; we would expect that it also underlies implicit biases. This is exactly what research in our lab has found: the historical proportion of enslaved populations at the county and state level predicts implicit bias today.[33] Places that relied on Black slave labor before abolition exhibit today higher pro-white bias among the white residents and lower pro-white bias among the Black residents. This effect persists even after controlling for self-reported attitudes. As we would predict from the bias of crowds theory, the relationship between the proportion of enslaved populations and implicit bias was mediated by structural inequalities like the proportion of Black people and white people in poverty, residential segregation, and intergenerational mobility of Black people and white people. Slavery and the ensuing generations of racial segregation and economic deprivation build the statistical regularities of racial inequality into the context. Chronic exposure to these structural inequalities maintains and exacerbates implicit bias.

Unfortunately, implicit biases are not merely cognitive reflections of our environment. Rather, they are influential aspects of our cognitive processes that change our behavior. The bias of crowds model suggests a recursive process such that inequalities of the past create the conditions for implicit biases to develop; and when they do, implicit biases contribute to the perpetuation of inequalities going forward. Said another way, implicit bias may be understood as both a cause and an effect of racial inequality.

There are many studies demonstrating that regional differences in implicit bias are associated with an increase in behaviors and outcomes that reinforce racial disparities. Such effects begin before children are born. An analysis of data from thirty-one million births across the United States found that the white-Black disparity in low birth weight is 14 percent higher in counties with high implicit bias.[34] During the SARS-CoV-2 global pandemic, anti-Black implicit biases of the white population across 957 counties predicted higher white and Black incidence of COVID-19 infection and a larger Black-white infection rate gap.[35] These are just specific instances of the larger pattern of racial health disparities following from geographic differences in implicit bias. A systematic review of the literature found evidence that all-cause mortality, cause-specific mortality, birth outcomes, cardiovascular outcomes, mental health, and self-rated health of racially minoritized groups are adversely affected by implicit biases.[36]

Implicit biases are also associated with the experiences of children. Counties with high levels of implicit bias show higher Black-white disparities in disciplinary suspension rates, and counties with higher levels of implicit bias among educators showed higher Black-white disparities in test scores and suspensions (after adjusting for several county-level covariates).[37] Similarly, county-level rates of anti-Black bias predict Black-white disparities in K–12 enrollment in gifted and talented programs such that high levels of bias predict large gaps and low levels of bias predict no gap.[38] U.S. states with higher levels of anti-Black implicit bias are also more likely to have lower adoption rates of Black foster children.[39]

Finally, several studies have demonstrated that regional implicit bias influences policing policy and behavior. Counties with higher anti-Black implicit bias have greater racial disparities in traffic stops.[40] Data from over two million residents across the United States also found that implicit biases predict more police officer use of lethal force against Black residents relative to the base rate of Black people in the population.[41] Researchers have also linked implicit bias to the problem of police militarization: regional differences in prejudice (including implicit bias) predict greater tax allocations for purchasing militarized police equipment.[42]

More research is needed to disentangle the many related factors involved in explaining racial disparities. Much of this work is relatively new and still developing. There is also some apparent overlap between the research linking historical events and policies to implicit biases and structural inequality. As we have recently argued, future researchers need to consider novel ways of incorporating these different factors into a coherent theoretical and statistical model.[43] Doing so will require collaboration between scholars from different fields like sociology, history, and policy in order to better account for the role of historical events, structural inequalities, and policy regimes, in addition to the usual set of predictors (implicit bias, explicit bias, and demographics) and policing, educational, and economic outcomes.

To address the massive public policy problems of racism and racial inequality in policing, education, economics, and health, racial justice advocates have turned to implicit bias as a focal point for intervention. Generally, this has taken the form of implicit bias trainings, whereby participants engaged in activities – such as perspective-taking, considering counter-stereotypical exemplars, meditating, or viewing empathy-building stimuli – designed to reduce implicit biases.[44] Unfortunately, meta-analyses and large-scale replications of such interventions have demonstrated that while they can successfully reduce implicit biases in the immediate time after intervention, they rarely have a sustained effect on implicit bias.[45] From a bias of crowds perspective, it is unsurprising that such interventions do not have lasting effects. If implicit bias is an emergent property of racial inequality in the social context, interventions that do not change the social context should leave implicit biases relatively unchanged.

By recognizing that context shapes implicit bias and behavior, researchers, policymakers, and practitioners can consider changing the context to reduce implicit bias.[46] At the highest level are societal-scale interventions that would redress historical and current inequality and thus radically change the context. Economist Ellora Derenoncourt and colleagues used economic data from 1860 to 2020 to simulate how economic conditions and policies influence the Black-white wealth gap.[47] Their analyses reveal that different combinations of policies that increase stock (such as lump sum reparations) and flow (such as facilitating financial diversification, stock equity, financial literacy, saving behaviors, and improving educational and labor market outcomes) in the Black community are feasible mechanisms for reducing wealth inequality over the coming decades.

On a smaller scale, individual organizations can reduce implicit bias by shifting organizational policies. Rather than having counter-stereotypical examples embedded in implicit bias training materials, organizations can work to have more counter-stereotypical minoritized group members in their ranks. Having an inclusive, equitable, and diverse team is a way to counteract the maintenance of negative intergroup biases.[48] This approach requires that organizations contend with biases in the hiring process that may hinder the hiring potential of racialized minority group members. To reduce the influence of implicit biases, decision-making processes can be predetermined and specified using hiring rubrics.[49] Though such rubrics can improve the hiring process, they need to follow evidence-based implementation to avoid perpetuating bias.[50] Another approach is to build in monitoring and accountability in hiring practices. Decision-makers who are held accountable for evaluating job candidates tend to show less pro-white biases.[51] Finally, people are more likely to be influenced by implicit biases when they are rushed, tired, distracted, or over-worked.[52] In a study of more than one thousand three hundred field experiments in classrooms, researchers found that discrimination rates against students from ethnic minority backgrounds were much lower when teachers were provided more time and resources in the classroom.[53]

Shifts in governmental, social, and workplace policies may be more challenging to implement compared with providing an implicit bias training, but policy changes may address the roots of the problem in ways that simple trainings cannot. Generations of public and organizational policy decisions resulted in the racial inequality we have today; the evidence suggests that we need equitable policies to counteract those effects.

Historically, the study of racism in psychological research has largely focused on interpersonal racism and has generally construed racism as an aspect of individual psychology, while neglecting the historical and structural aspects of racism.[54] The bias of crowds model is a theoretical framework that explains why the modern shift toward racial egalitarianism in attitudes has not resulted in diminished racial inequality. It also accounts for the many research findings that are inconsistent with the perspective that implicit bias is a stable aspect of individual psychology.

The other benefit of the bias of crowds model is that it makes efficient use of existing data and theory. Research that links together policy, structural inequality, and psychological measures has been limited by the availability of geo-coded "big data" on these topics. The recent explosion of research linking widescale policies like redlining to health outcomes (for example) is in no small part due to the increased availability of such data. Analyses using these data reveal more evidence that the bias of crowds model is a social psychological model consistent with the sociological theory of systemic racism. Ultimately, the model connects many forms of racism – structural, systemic, implicit, explicit, cultural, historical, current – under one testable theoretical perspective.

Finally, the bias of crowds model reinforces what many in sociology, economics, history, and policy have articulated in their work: we need to consider systems to understand and ameliorate racial inequality. The bias of crowds model shifts the focus of research designed to address inequality to consider the impact of changing the broader social context.

ABOUT THE AUTHORS

Manuel J. Galvan is a PhD candidate in social psychology at the University of North Carolina at Chapel Hill. He is the Cofounder and Cohost of the *A Bit More Complicated* science podcast. He has published in such journals as *Science*, *Annals of Behavioral Medicine*, and *American Psychologist*.

B. Keith Payne is Professor of Psychology and Neuroscience at the University of North Carolina at Chapel Hill. He is the author of *The Broken Ladder: How Inequality Affects the Way We Think, Live, and Die* (2017) and has recently published in such journals as *Psychological Science*, *Journal of Experimental Psychology*, and *Proceedings of the National Academy of Sciences*.

ENDNOTES

[1] Eric Holloway, "Yes, ChatGPT Is Sentient–Because It's Really Humans in the Loop," *Mind Matters*, December 26, 2022, https://mindmatters.ai/2022/12/yes-chatgpt-is-sentient-because-its-really-humans-in-the-loop.

[2] Tony H. Tran, "OpenAI's Impressive New Chatbot Isn't Immune to Racism," *The Daily Beast*, December 6, 2022, https://www.thedailybeast.com/openais-impressive-chatgpt-chatbot-is-not-immune-to-racism.

[3] Sarah Patton Moberg, Maria Krysan, and Deanna Christianson, "Racial Attitudes in America," *Public Opinion Quarterly* 83 (2) (2019): 450–471, https://doi.org/10.1093/poq/nfz014.

[4] Juliana M. Horowitz, Anna Brown, and Kiana Cox, "Race in America 2019," Pew Research Center, April 9, 2019, https://www.pewsocialtrends.org/wp-content/uploads/sites/3/2019/04/PewResearchCenter_RaceStudy_FINAL-1.pdf.

[5] Economic Research Service, *Food Security Status of U.S. Households in 2021* (Washington, D.C.: Economic Research Service, 2023), https://www.ers.usda.gov/topics/food-nutrition-assistance/food-security-in-the-u-s/key-statistics-graphics/#householdtype; Jonnelle Marte, "Gap in U.S. Black and White Unemployment Rates Is Widest in Five Years," Reuters, July 2, 2020, https://www.reuters.com/article/us-usa-economy-unemployment-race/gap-in-u-s-black-and-white-unemployment-rates-is-widest-in-five-years-idUSKBN2431X7; and U.S. Census Bureau, "Historical Poverty Tables: People and Families–1959 to 2021," last modified January 30, 2023, https://www.census.gov/data/tables/time-series/demo/income-poverty/historical-poverty-people.html.

[6] Lisa Dettling, Joanne Hsu, Lindsay Jacobs, Kevin B. Moore, and Jeffrey Thompson, "Recent Trends in Wealth-Holding by Race and Ethnicity: Evidence from the Survey of Consumer Finances," FEDS Notes, September 27, 2017, https://www.federalreserve.gov/econres/notes/feds-notes/recent-trends-in-wealth-holding-by-race-and-ethnicity-evidence-from-the-survey-of-consumer-finances-20170927.html; and Valerie Wilson and Jhacova Williams, "Racial and Ethnic Income Gaps Persist Amid Uneven Growth in Household Incomes," Economic Policy Institute Working Economics Blog, September 11, 2019, https://www.epi.org/blog/racial-and-ethnic-income-gaps-persist-amid-uneven-growth-in-household-incomes.

[7] Katrin Auspurg, Andreas Schneck, and Thomas Hinz, "Closed Doors Everywhere? A Meta-Analysis of Field Experiments on Ethnic Discrimination in Rental Housing Mar-

kets," *Journal of Ethnic and Migration Studies* 45 (1) (2019): 95–114, https://doi.org/10.1080/1369183X.2018.1489223; and Lincoln Quillian, Devah Pager, Ole Hexel, and Arnfinn H. Midtbøen, "Meta-Analysis of Field Experiments Shows No Change in Racial Discrimination in Hiring over Time," *Proceedings of the National Academy of Sciences* 114 (41) (2017): 10870–10875, https://doi.org/10.1073/pnas.1706255114.

[8] Joe Feagin, *Systemic Racism: A Theory of Oppression* (Abingdon-on-Thames: Routledge, 2013); and Joe Feagin and Kimberley Ducey, *Racist America: Roots, Current Realities, and Future Reparations* (Abingdon-on-Thames: Taylor & Francis, 2000).

[9] Brian A. Nosek, Frederick L. Smyth, Jeffrey J. Hansen, et al., "Pervasiveness and Correlates of Implicit Attitudes and Stereotypes," *European Review of Social Psychology* 18 (1) (2007): 36–88, https://doi.org/10.1080/10463280701489053; and Kate A. Ratliff, Jennifer Howell, Colin Smith, et al., "Documenting Bias from 2007–2015," unpublished manuscript, last updated February 28, 2020, https://osf.io/rfzhu (accessed November 16, 2023).

[10] Anthony G. Greenwald, Debbie E. McGhee, and Jordan L. K. Schwartz, "Measuring Individual Differences in Implicit Cognition: The Implicit Association Test," *Journal of Personality and Social Psychology* 74 (6) (1998): 1464–1480, https://doi.org/10.1037/0022-3514.74.6.1464.

[11] Richard E. Petty, Zakary L. Tormala, Pablo Brinol, and W. Blair G. Jarvis, "Implicit Ambivalence from Attitude Change: An Exploration of the PAST Model," *Journal of Personality and Social Psychology* 90 (1) (2006): 21–41, https://doi.org/10.1037/0022-3514.90.1.21; Laurie A. Rudman, "Sources of Implicit Attitudes," *Current Directions in Psychological Science* 13 (2) (2004): 79–82; and Timothy D. Wilson, Samuel Lindsey, and Tonya Y. Schooler, "A Model of Dual Attitudes," *Psychological Review* 107 (1) (2000): 101–126, https://doi.org/10.1037/0033-295X.107.1.101.

[12] Greenwald, McGhee, and Schwartz, "Measuring Individual Differences in Implicit Cognition"; and B. Keith Payne, "Conceptualizing Control in Social Cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping," *Journal of Personality and Social Psychology* 89 (4) (2005): 488–503, https://doi.org/10.1037/0022-3514.89.4.488.

[13] Erin Cooley and B. Keith Payne, "Using Groups to Measure Intergroup Prejudice," *Personality and Social Psychology Bulletin* 43 (1) (2017): 46–59, https://doi.org/10.1177/0146167216675331; Bertram Gawronski, Mike Morrison, Curtis E. Phills, and Silvia Galdi, "Temporal Stability of Implicit and Explicit Measures: A Longitudinal Analysis," *Personality and Social Psychology Bulletin* 43 (3) (2017): 300–312, https://doi.org/10.1177/0146167216684131; Wilhelm Hofmann, Jan De Houwer, Marco Perugini, Frank Baeyens, and Geert Crombez, "Evaluative Conditioning in Humans: A Meta-Analysis," *Psychological Bulletin* 136 (3) (2010): 390–421, https://pubmed.ncbi.nlm.nih.gov/20438144; and Keith Payne and Kristjen Lundberg, "The Affect Misattribution Procedure: Ten Years of Evidence on Reliability, Validity, and Mechanisms," *Social and Personality Psychology Compass* 8 (12) (2014): 672–686, https://doi.org/10.1111/spc3.12148.

[14] Andrew Scott Baron and Mahzarin R. Banaji, "The Development of Implicit Attitudes: Evidence of Race Evaluations from Ages 6 and 10 and Adulthood," *Psychological Science* 17 (1) (2006): 53–58, https://doi.org/10.1111/j.1467-9280.2005.01664.x; Yarrow Dunham, Eva E. Chen, and Mahzarin R. Banaji, "Two Signatures of Implicit Intergroup Attitudes: Developmental Invariance and Early Enculturation," *Psychological Science* 24 (6) (2013): 860–868, https://doi.org/10.1177/0956797612463081; and Brandon D. Stewart,

William Von Hippel, and Gabriel A. Radvansky, "Age, Race, and Implicit Prejudice: Using Process Dissociation to Separate the Underlying Components," *Psychological Science* 20 (2) (2009): 164–168, https://doi.org/10.1111/j.1467-9280.2009.02274.x.

[15] Anthony G. Greenwald, T. Andrew Poehlman, Eric Luis Uhlmann, and Mahzarin R. Banaji, "Understanding and Using the Implicit Association Test: III. Meta-Analysis of Predictive Validity," *Journal of Personality and Social Psychology* 97 (1) (2009): 17–41, https://doi.org/10.1037/a0015575; and Frederick L. Oswald, Gregory Mitchell, Hart Blanton, et al., "Predicting Ethnic and Racial Discrimination: A Meta-Analysis of IAT Criterion Studies," *Journal of Personality and Social Psychology* 105 (2) (2013): 171–192.

[16] Irene V. Blair and Elizabeth Brondolo, "Moving Beyond the Individual: Community-Level Prejudice and Health," *Social Science & Medicine* 183 (2017): 169–172, https://doi.org/10.1016/j.socscimed.2017.04.041; Eric Hehman, Jessica K. Flake, and Jimmy Calanchini, "Disproportionate Use of Lethal Force in Policing Is Associated With Regional Racial Biases of Residents," *Social Psychological and Personality Science* 9 (4) (2018): 393–401, https://doi.org/10.1177/1948550617711229; James R. Rae, Anna Kaisa Newheiser, and Kristina R. Olson, "Exposure to Racial Out-Groups and Implicit Race Bias in the United States," *Social Psychological and Personality Science* 6 (5) (2015): 535–543; and Brian A. Nosek, Frederick L. Smyth, Natarajan Sriram, et al., "National Differences in Gender-Science Stereotypes Predict National Sex Differences in Science and Math Achievement," *Proceedings of the National Academy of Sciences* 106 (26) (2009): 10593–10597, https://doi.org/10.1073/pnas.0809921106.

[17] B. Keith Payne, Heidi A. Vuletich, and Kristjen B. Lundberg, "The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice," *Psychological Inquiry* 28 (4) (2017): 233–248, https://doi.org/10.1080/1047840X.2017.1335568.

[18] Peter J. Rentfrow, Samuel D. Gosling, Markus Jokela, et al., "Divided We Stand: Three Psychological Regions of the United States and Their Political, Economic, Social, and Health Correlates," *Journal of Personality and Social Psychology* 105 (6) (2013): 996–1012, https://doi.org/10.1037/a0034434.

[19] Peter J. Rentfrow, Markus Jokela, and Michael E. Lamb, "Regional Personality Differences in Great Britain," *PLOS ONE* 10 (3) (2015), https://doi.org/10.1371/journal.pone.0122245.

[20] Thomas J. Allen, Jeffrey W. Sherman, and Karl Christoph Klauer, "Social Context and the Self-Regulation of Implicit Bias," *Group Processes and Intergroup Relations* 13 (2) (2010): 137–149, https://doi.org/10.1177/1368430209353635; Jamie Barden, William W. Maddux, Richard E. Petty, and Marilynn B. Brewer, "Contextual Moderation of Racial Bias: The Impact of Social Roles on Controlled and Automatically Activated Attitudes," *Journal of Personality and Social Psychology* 87 (1) (2004): 5–22, https://doi.org/10.1037/0022-3514.87.1.5; and Brian S. Lowery, Curtis D. Hardin, and Stacey Sinclair, "Social Influence Effects on Automatic Racial Prejudice," *Journal of Personality and Social Psychology* 81 (5) (2001): 842–855, https://doi.org/10.1037//0022-3514.81.5.842.

[21] Barden, Maddux, Petty, and Brewer, "Contextual Moderation of Racial Bias"; and Bernd Wittenbrink, Charles M. Judd, and Bernadette Park, "Spontaneous Prejudice in Context: Variability in Automatically Activated Attitudes," *Journal of Personality and Social Psychology* 81 (5) (2001): 815–827, http://doi.org/10.1037/0022-3514.81.5.815.

[22] Andrea R. French, Timothy M. Franz, Laura L. Phelan, and Bruce E. Blaine, "Reducing Muslim/Arab Stereotypes through Evaluative Conditioning," *The Journal of Social Psychology* 153 (1) (2013): 6–9, https://doi.org/10.1080/00224545.2012.706242.

[23] Hofmann, De Houwer, Perugini, et al., "Evaluative Conditioning in Humans."

[24] Wittenbrink, Judd, and Park, "Spontaneous Prejudice in Context"; and Nilanjana Dasgupta and Anthony G. Greenwald, "On the Malleability of Automatic Attitudes: Combating Automatic Prejudice with Images of Admired and Disliked Individuals," *Journal of Personality and Social Psychology* 81 (5) (2001): 800–814, https://doi.org/10.1037/0022-3514.81.5.800.

[25] Mary Mederios Kent, "Immigration and America's Black Population," *Population Bulletin* 62 (4) (2007), https://u.demog.berkeley.edu/~jrw/Biblio/Eprints/PRB/files/62.4immigration.pdf.

[26] Samantha Longman-Mills, Carole Mitchell, and Wendel Abel, "The Psychological Trauma of Slavery: The Jamaican Case Study," *Social and Economic Studies* 68 (3–4) (2019): 79–102.

[27] Avidit Acharya, Matthew Blackwell, and Maya Sen, "The Political Legacy of American Slavery," *The Journal of Politics* 78 (3) (2016): 621–641; and Nathan Nunn and Leonard Wantchekon, "The Slave Trade and the Origins of Mistrust in Africa," *American Economic Review* 101 (7) (2011): 3221–3252, https://doi.org/10.1257/aer.101.7.3221.

[28] Ellora Derenoncourt, Chi Hyun Kim, Moritz Kuhn, and Moritz Schularick, "Wealth of Two Nations: The U.S. Racial Wealth Gap, 1860–2020," NBER Working Paper 30101 (Cambridge, Mass.: National Bureau of Economic Research, 2022).

[29] Nancy Krieger, Jarvis T. Chen, Brent A. Coull, et al., "Jim Crow and Premature Mortality among the U.S. Black and White Population, 1960–2009: An Age-Period-Cohort Analysis," *Epidemiology* 25 (4) (2014): 494–504, https://doi.org/10.1097/EDE.0000000000000104.

[30] Jhacova Williams, "Historical Lynchings and the Contemporary Voting Behavior of Blacks," *American Economic Journal: Applied Economics* 14 (3) (2022): 224–253, https://doi.org/10.1257/app.20190549.

[31] Heather A. O'Connell, "Historical Shadows: The Links Between Sundown Towns and Contemporary Black–White Inequality," *Sociology of Race and Ethnicity* 5 (3) (2019): 311–325, https://doi.org/10.1177/2332649218761979.

[32] Sonali Bose, Jaime Madrigano, and Nadia N. Hansel, "When Health Disparities Hit Home: Redlining Practices, Air Pollution, and Asthma," *American Journal of Respiratory and Critical Care Medicine* 206 (7) (2022): 803–804, https://doi.org/10.1164/rccm.202206-1063ed; Clese E. Erikson, Randl B. Dent, Yoon Hong Park, and Qian Luo, "Historic Redlining and Contemporary Behavioral Health Workforce Disparities," *JAMA Network Open* 5 (4) (2022), https://doi.org/10.1001/jamanetworkopen.2022.9494; Nick Graetz and Michael Esposito, "Historical Redlining and Contemporary Racial Disparities in Neighborhood Life Expectancy," *Social Forces* 102 (1) (2023): 1–22, https://doi.org/10.1093/sf/soac114; Min Li and Faxi Yuan, "Historical Redlining and Resident Exposure to COVID-19: A Study of New York City," *Race and Social Problems* 14 (2) (2022): 85–100, https://doi.org/10.1007%2Fs12552-021-09338-z; Min Li and Faxi Yuan, "Historical Redlining and Food Environments: A Study of 102 Urban Areas in the United States," *Health & Place* 75 (2022): 102775, https://doi.org/10.1016/j.healthplace.2022.102775; Jeffrey Mitchell and Guilherme Kenji Chihaya, "Tract Level Associations Between Historical Residential Redlining and Contemporary Fatal Encounters with Police," *Social Science & Medicine* 302 (2022): 114989, https://doi.org/10.1016/j.socscimed.2022.114989; Issam Motairek, Zhuo Chen, Mohamed H. E. Makhlouf, et al., "Historical Neighbour-

hood Redlining and Contemporary Environmental Racism," *Local Environment* 28 (4) (2023): 518–528; Issam Motairek, Eun Kyung Lee, Scott Janus, et al., "Historical Neighborhood Redlining and Contemporary Cardiometabolic Risk," *Journal of the American College of Cardiology* 80 (2) (2022): 171–175, https://doi.org/10.1016/j.jacc.2022.05.010; Richard Casey Sadler, Thomas Walter Wojciechowski, Pamela Trangenstein, et al., "Linking Historical Discriminatory Housing Patterns to the Contemporary Alcohol Environment," *Applied Spatial Analysis and Policy* 16 (2) (2023): 561–581, https://doi.org/10.1007/s12061-022-09493-9; Leah H. Schinasi, Chahita Kanungo, Zachary Christman, et al., "Associations between Historical Redlining and Present-Day Heat Vulnerability Housing and Land Cover Characteristics in Philadelphia, PA," *Journal of Urban Health* 99 (1) (2022): 134–145, https://doi.org/10.1007/s11524-021-00602-6; Elli Schwartz, Nathaniel Onnen, Peter F. Craigmile, and Megan E. Roberts, "The Legacy of Redlining: Associations between Historical Neighborhood Mapping and Contemporary Tobacco Retailer Density in Ohio," *Health & Place* 68 (2021): 102529, https://doi.org/10.1016/j.healthplace.2021.102529; and Mudia Uzzi, Kyle T. Aune, Lea Marineau, et al., "An Intersectional Analysis of Historical and Contemporary Structural Racism on Non-Fatal Shootings in Baltimore, Maryland," *Injury Prevention* 29 (1) (2023): 85–90, http://doi.org/10.1136/ip-2022-044700.

33 B. Keith Payne, Heidi A. Vuletich, and Jazmin L. Brown-Iannuzzi, "Historical Roots of Implicit Bias in Slavery," *Proceedings of the National Academy of Sciences* 116 (24) (2019): 11693–11698, https://doi.org/10.1073/pnas.1818816116.

34 Jacob Orchard and Joseph Price, "County-Level Racial Prejudice and the Black-White Gap in Infant Health Outcomes," *Social Science & Medicine* 100 (181) (2017): 191–198, https://doi.org/10.1016/j.socscimed.2017.03.036.

35 Marilyn D. Thomas, Eli K. Michaels, Sean Darling-Hammond, et al., "Whites' County-Level Racial Bias, COVID-19 Rates, and Racial Inequities in the United States," *International Journal of Environmental Research and Public Health* 17 (22) (2020): E8695, https://doi.org/10.3390/ijerph17228695.

36 Eli K. Michaels, Christine Board, Mahasin S. Mujahid, et al., "Area-Level Racial Prejudice and Health: A Systematic Review," *Health Psychology* 41 (3) (2022): 211–224, https://doi.org/10.1037%2Fhea0001141; and Tessa E. S. Charlesworth and Mahzarin R. Banaji, "Evidence of Covariation Between Regional Implicit Bias and Socially Significant Outcomes in Healthcare, Education, and Law Enforcement," in *Handbook on Economics of Discrimination and Affirmative Action*, ed. Ashwini Deshpande (Berlin: Springer Nature, 2023), 593–613.

37 Mark J. Chin, David M. Quinn, Tasminda K. Dhaliwal, and Virginia S. Lovison, "Bias in the Air: A Nationwide Exploration of Teachers' Implicit Racial Attitudes, Aggregate Bias, and Student Outcomes," *Educational Researcher* 49 (8) (2020): 566–578, https://doi.org/10.3102/0013189X20937240.

38 Francis A. Pearman and Ebony O. McGee, "Anti-Blackness and Racial Disproportionality in Gifted Education," *Exceptional Children* 88 (4) (2022): 359–380.

39 Sarah Beth Bell, Rachel Farr, Eugene Ofosu, Eric Hehman, and C. Nathan DeWall, "Implicit Bias Predicts Less Willingness and Less Frequent Adoption of Black Children More Than Explicit Bias," *The Journal of Social Psychology* 163 (4) (2023): 554–565.

40 Pierce D. Ekstrom, Joel M. Le Forestier, and Calvin K. Lai, "Racial Demographics Explain the Link between Racial Disparities in Traffic Stops and County-Level Racial Attitudes," *Psychological Science* 33 (4) (2022): 497–509, https://doi.org/10.1177/09567976211053573;

and Marleen Stelter, Iniobong Essien, Carsten Sander, and Juliane Degner, "Racial Bias in Police Traffic Stops: White Residents' County-Level Prejudice and Stereotypes Are Related to Disproportionate Stopping of Black Drivers," *Psychological Science* 33 (4) (2022): 483–496, https://doi.org/10.1177/09567976211051272.

[41] Hehman, Flake, and Calanchini, "Disproportionate Use of Lethal Force in Policing Is Associated With Regional Racial Biases of Residents."

[42] Tyler Jimenez, Peter J. Helm, and Jamie Arndt, "Racial Prejudice Predicts Police Militarization," *Psychological Science* 33 (12) (2022): 2009–2026, https://doi.org/10.1177/09567976221112936.

[43] B. Keith Payne and Julian M. Rucker, "Explaining the Spatial Patterning of Racial Disparities in Traffic Stops Requires a Structural Perspective: Further Reflections on Stelter et al. (2022) and Ekstrom et al. (2022)," *Psychological Science* 33 (4) (2022): 666–668, https://doi.org/10.1177/09567976211056641.

[44] Mimi V. Chapman, William J. Hall, Kent Lee, et al., "Making a Difference in Medical Trainees' Attitudes toward Latino Patients: A Pilot Study of an Intervention to Modify Implicit and Explicit Attitudes," *Social Science & Medicine* 199 (2018): 202–208, https://doi.org/10.1016%2Fj.socscimed.2017.05.013; Dasgupta and Greenwald, "On the Malleability of Automatic Attitudes"; Alexander J. Stell and Tom Farsides, "Brief Loving-Kindness Meditation Reduces Racial Bias, Mediated by Positive Other-Regarding Emotions," *Motivation and Emotion* 40 (1) (2016): 140–147, https://doi.org/10.1007/s11031-015-9514-x; and Andrew R. Todd, Galen V. Bodenhausen, Jennifer A. Richeson, and Adam D. Galinsky, "Perspective Taking Combats Automatic Expressions of Racial Bias," *Journal of Personality and Social Psychology* 100 (6) (2011): 1027–1042.

[45] Patrick S. Forscher, Calvin K. Lai, Jordan R. Axt, et al., "A Meta-Analysis of Procedures to Change Implicit Measures," *Journal of Personality and Social Psychology* 117 (3) (2019): 522–559, https://doi.org/10.1037/pspa0000160; and Calvin K. Lai, Allison L. Skinner, Erin Cooley, et al., "Reducing Implicit Racial Preferences: II. Intervention Effectiveness Across Time," *Journal of Experimental Psychology: General* 145 (8) (2016): 1001–1016, https://doi.org/10.1037/xge0000179.

[46] B. Keith Payne and Heidi A. Vuletich, "Policy Insights from Advances in Implicit Bias Research," *Policy Insights from the Behavioral and Brain Sciences* 5 (1) (2018): 49–56, https://doi.org/10.1177/2372732217746190.

[47] Derenoncourt, Kim, Kuhn, and Schularick, *Wealth of Two Nations*.

[48] George B. Cunningham, "The Influence of Group Diversity on Intergroup Bias following Recategorization," *The Journal of Social Psychology* 146 (5) (2006): 533–547.

[49] Eric Luis Uhlmann and Geoffrey L. Cohen, "Constructed Criteria: Redefining Merit to Justify Discrimination," *Psychological Science* 16 (6) (2005): 474–480.

[50] Dawn Culpepper, Damani White-Lewis, KerryAnn O'Meara, Lindsey Templeton, and Julia Anderson, "Do Rubrics Live Up to Their Promise? Examining How Rubrics Mitigate Bias in Faculty Hiring," *The Journal of Higher Education* 94 (7) (2023): 823–850, https://doi.org/10.1080/00221546.2023.2168411.

[51] Thomas E. Ford, Frank Gambino, Hanjoon Lee, et al., "The Role of Accountability in Suppressing Managers' Preinterview Bias against African-American Sales Job Applicants," *Journal of Personal Selling & Sales Management* 24 (2) (2004): 113–124.

[52] Payne, "Conceptualizing Control in Social Cognition"; B. Keith Payne, "Weapon Bias: Split-Second Decisions and Unintended Stereotyping," *Current Directions in Psychological Science* 15 (6) (2006): 287–291, https://doi.org/10.1111/j.1467-8721.2006.00454.x; and Jeffrey W. Sherman, Bertram Gawronski, Karen Gonsalkorale, et al., "The Self-Regulation of Automatic Associations and Behavioral Impulses," *Psychological Review* 115 (2) (2008): 314–335, https://doi.org/10.1037/0033-295x.115.2.314.

[53] Simon Calmar Andersen and Thorbjørn Sejr Guul, "Reducing Minority Discrimination at the Front Line–Combined Survey and Field Experimental Evidence," *Journal of Public Administration Research and Theory* 29 (3) (2019): 429–444, https://doi.org/10.1093/jopart/muy083.

[54] Sophie Trawalter, D-J Bart-Plange, and Kelly M. Hoffman, "A Socioecological Psychology of Racism: Making Structures and History More Visible," *Current Opinion in Psychology* 32 (2020): 47–51, https://doi.org/10.1016/j.copsyc.2019.06.029.