

This is a section of [doi:10.7551/mitpress/14922.001.0001](https://doi.org/10.7551/mitpress/14922.001.0001)

Open Minded

Searching for Truth about the Unconscious Mind

By: Ben R. Newell, David R. Shanks

Citation:

Open Minded: Searching for Truth about the Unconscious Mind

By: Ben R. Newell, David R. Shanks

DOI: 10.7551/mitpress/14922.001.0001

ISBN (electronic): 9780262375375

Publisher: The MIT Press

Published: 2023

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open



The MIT Press

5 Rethinking Unconscious Bias

The world is a very unequal place. In most Western companies and institutions, women and individuals from ethnic minorities are paid less and are less likely to be promoted than men. Of the top five hundred companies in the United States, only 41 had female chief executives in 2021.¹ Similar inequalities can be seen almost everywhere one looks. In Australia around forty people per 10,000 of the male population are in prison, but the figure for those of indigenous origin is a staggering 438.²

Bias refers to irrational preferences or tendencies that favor some choices over others.³ Central to the concept is the idea that someone manifesting a bias tends to be impervious to contrary evidence or remediation. If the selection panel for hiring new members of a prestigious orchestra continues to prefer male over female musicians despite objective evidence that the candidates have equal musical ability, no one would disagree that this is an irrational and undesirable bias as well as being unfair.

Males and females are of course roughly equal in the population overall. But few would argue that because more men than women are incarcerated, the prison system is therefore biased against men. Rather the prison population reflects the overall difference in criminality rates between the sexes. The point is that a bias is a bias only if it runs counter to relevant hard evidence and impartiality. Gordon Allport, one of the pioneers of social psychology, coined the phrase “thinking ill of others without sufficient warrant” to capture the idea that misalignment with objective facts is central to concepts such as bias and prejudice.

Establishing sufficient or insufficient warrant can quickly become quite complex. The University of Oxford has been frequently castigated because so few of its students come from ethnic minorities. But is this a bias in the true

sense of the term? About 17 percent of ethnic minority students are offered a place in contrast to about 26 percent of White applicants. Yes, it's a bias.

But hang on. Courses vary in how competitive they are, and it turns out that Black students are more likely than White ones to apply for the most competitive courses at Oxford (29 percent of Black students apply for medicine, a highly competitive subject, against only 7 percent of White applicants). So the fact that relatively fewer Black students are admitted to the university could be reflective of a bias, or instead could be a consequence of their tendency to apply for highly selective courses combined with a color-blind and unbiased admissions system.

But hang on. If we look at specific courses such as medicine, we find that highly able students (ones who got top grades in their final school exams) are much more likely to be offered a place if they're White (43 percent) than if they're Black (22 percent). So it is a bias after all.⁴ The point is that establishing whether a preference is a bias or not—whether it is warranted—can be very difficult when the data on which the preference should legitimately be based are hard to quantify and unpack.

When it comes to examining the behavior of individuals rather than systems, establishing bias is just as difficult. Proving that people show unconscious or implicit biases, preferences that they're not even aware of, is harder still. Yet this notion has gained a great deal of traction in recent years, especially among employers, who spend vast sums annually on diversity training programs and interventions designed to eliminate prejudice in the workplace. It is now commonplace for individuals sitting on selection panels to be required to undergo unconscious bias training, offered with the entirely laudable aim of helping us to recognize bias and intolerance in ourselves and others and alerting us to conditions and circumstances in which we may be unintentionally susceptible to bias. Training programs designed to suppress unconscious bias use a range of techniques, such as encouraging us to deliberate rather than make snap judgments, reconsider reasons for or against our decisions, question cultural stereotypes, and monitor each other.⁵ Yet the evidence that diversity training programs are effective at changing workplace behavior is murky at best,⁶ and more worrying, scant regard has been given to the possibility that such programs could actually be harmful. There is research clearly showing, for instance, that some interventions designed to reduce prejudice—particularly ones that emphasize societal norms against prejudice—may backfire and in fact increase it.⁷ It

is therefore imperative that the case for unconscious bias is examined critically. This is the subject of this chapter.

A Kernel of Truth?

There can be little doubt that the human mind is capable of making extraordinarily subtle snap judgments and that our decisions are often based on only some of the relevant evidence. For example, a brief glimpse of around a tenth of a second is sufficient for us to decide whether a face is trustworthy or competent.⁸ If such judgments are often inaccurate, it would clearly be irrational to base our decisions and preferences on such snap judgments. It is not a good policy to trust another person simply because something in her face makes us feel intuitively that she is trustworthy. But research findings illustrate convincingly that these snap judgments often tend to have some validity.

In one study, the cooperativeness of rural Senegalese men and women was established by asking them to play a four-person public goods game.⁹ In this game, each player was initially given 200 grams of rice, which they could divide between a private and a public endowment. The rules of the game dictated that the total amount contributed to the public endowment by all four players was doubled and equally shared. Thus, players came away with an equal share of the public endowment plus whatever they held back as their private endowment. If each player allocated all of their rice to the private pool, they all ended the game with the same amount of rice as they started: 200 grams. If they all allocated all of their rice to the public endowment, each finished the game with 400 grams. However, it is possible to free-ride: a player who contributes nothing to the pool nonetheless benefits when the enlarged pool is shared. The game therefore measures the extent to which each player is willing to cooperate in a public endeavor to pool and multiply a good, and some of the Senegalese playing the game were found to be cooperative and others selfish. Can each person's cooperativeness be judged from his or her face? In the next phase of the experiment, pairs of photographs of these Senegalese players were shown to volunteers over two thousand miles away in Montpellier, France, who were asked to decide which member of each pair of faces was the more selfish person. Remarkably, the selfishness of male (though not female) players could be judged at levels better than pure guessing. A brief look at the face of a

male from a different culture is sufficient to extract some valid information about that person's cooperativeness.

One area in which bias has, understandably, been extensively explored is in relation to gender and employment. In one controlled study, a job application was created, purportedly submitted by a recent science graduate who was applying for a laboratory manager position.¹⁰ The application was assessed by a large number of science faculty at several US universities who rated it for how employable and competent they judged the applicant and what starting salary they would offer. Crucially, the applicant either had a female (Jennifer) or a male (John) name. Despite the fact that everything else in the application was identical, the female applicant was judged less competent and employable than the male one and offered a lower starting salary. Since the only difference between the applications was the individual's name and implied gender, the results seem to strongly imply gender bias in job evaluations and hiring decisions.

The problem with this conclusion is that discrimination and bias are not the same thing. Any employer who hires women at lower rates than men even when they have equivalent qualifications is guilty of discrimination in the legal sense of the term. It is contrary to employment law to base hiring decisions on gender. But this does not necessarily make it irrational to do so. There is much evidence showing that the lifetime productivity of female researchers is lower on average than that of males, not least because of career breaks to have children.¹¹ If only a minimal amount of evidence about a candidate is available, then gender is a valid signal of likely productivity. (Of course, one can legitimately question whether lifetime productivity is a fair criterion against which to measure success.) It is quite right that discrimination is illegal and that employers put policies in place to mitigate it, but that doesn't mean it's an unwarranted bias in the sense we're investigating here.

Evidence from studies like this is limited because they examine hypothetical hiring decisions involving rather generic job applications. Under such circumstances it could be rational to include base-rate information such as gender differences in productivity in one's decisions (whether it is moral or legal to do so is a very different matter). Consider this analogy. Imagine that you run a removal company and your employees have to be fit and strong enough to do a lot of lifting of furniture, boxes, pianos, and so on. You need to hire a new employee. If all you have to go on are job

applications that reveal contact details and mundane information about educational qualifications, then choosing only to interview male applicants would be perfectly sensible given that males on average are stronger than females. If, on the other hand, you receive applications from individuals, all of whom describe previous experience working for removal companies and excellent job appraisals and letters of recommendation, then it would be counterproductive to pay any attention to gender. The point is that the more generic and nonspecific an application, the more sensible it is to put weight on statistical signals such as gender. To reiterate, we are talking here about whether decisions are biased or impartial, not about whether they're legally or morally acceptable or fair.

Proving the point, when we look at more realistic hiring decisions, involving much more comprehensive information about the candidate, gender bias against women is much harder to detect. Much of this research has focused on hiring in university departments. In the largest investigation undertaken to date, Wendy Williams and Stephen Ceci asked faculty in a range of science disciplines to evaluate realistic job applicants for a tenure-track position.¹² The application materials were highly realistic and included detailed notes describing the search committee chair's evaluation of each candidate, with comments such as "Z's faculty job talk/interview score was 9.5/10. At dinner with the committee, she impressed everyone as being a confident and professional individual with a great deal to offer the department. During our private meeting, Z was enthusiastic about our department, and there did not appear to be any obstacles if we decided to offer her the job." Of course, the same materials were rated by some faculty under a female identity and by others under a male identity.

Remarkably, female applicants were favored over males by a ratio of about 2:1. This pattern held over both math-intensive fields like engineering as well as less math-intensive ones like biology, the only exception being economics, where no preference for either gender was found. Indeed, evidence on the actual hiring rates of men and women for faculty positions at US universities shows a similar preference for, not against, women (although fewer women apply overall). Between 2002 and 2004, for instance, around 20 percent of all applicants for faculty math positions in universities in the United States were women, but 32 percent of those offered jobs were women.¹³ Of course higher education is only one among many employment sectors, and there is a persistent gender pay gap in almost all types of

employment. But evidence for the idea that psychological factors such as implicit bias play a role in hiring decisions is distinctly underwhelming.¹⁴ To state the obvious, judgments and preferences can't reflect unconscious biases if they're not even biases.

Similar patterns emerge in other contexts in which gender (or race) is one among many indicators. For example, in scientific publishing, researchers or research teams send their often lengthy and detailed reports to journal editors who solicit evaluations from expert reviewers and decide whether to publish them. These are important decisions as promotion and pay in universities and other research organizations depend hugely on one's track record of scholarly publications. Although women make up less than one-third of the authors of all scientific publications globally, articles submitted to journals by female authors are no less likely to be published than those by male authors.¹⁵ When gender is rendered an irrelevant signal, it is ignored in decision making. It is irrelevant in this case because the decision about whether to publish an article should be (and is) based on its intrinsic quality.

Indeed far from supporting the common (mis)conception that biases rest on inaccurate stereotypes about social groups, the evidence is much more consistent with what has been dubbed the "kernel of truth" hypothesis, which maintains that many of our stereotypes tend to reflect reality, at least to a first approximation. If people generally hold the stereotype that women tend to be more conscientious but less extroverted than men, perhaps this is because it's actually the case—which indeed is what the research findings show. The social psychologist Lee Jussim has argued forcefully that the automatic assumption that stereotypes are inaccurate grossly misrepresents reality.¹⁶

Unconscious Associations

Alongside this sort of quasi-naturalistic research on implicit bias, there has been much investigation using laboratory tasks, most famously the well-known implicit association test—the IAT. This test for measuring unconscious attitudes, devised by Anthony Greenwald and his colleagues at the University of Washington, takes the form of a simple set of decisions. Imagine that you're shown a series of words one at a time on a computer screen, such as *lovely*, *dirty*, *friendship*, *rotten*, *appealing*, *evil*, *cheerful*, *failure*, and you have to press a left key for positive words (such as *cheerful*) and a right key

for negative words (such as *dirty*). The speed with which you make these decisions is recorded. Next you're shown photographs of Black or White people and again required to press one key for Black faces and the other key for White faces. This is all easy.

Now we get to the interesting part. You see a sequence of randomly interspersed words and faces and have to press the left key for negative words and Black faces and the right key for positive words and White faces. It should be intuitive that if you have an unconscious bias against Black people, it will be relatively easy to respond rapidly in this stage of the experiment as the separate decisions (Black—left; White—right; negative—left; positive—right) are compatible in the sense that the left key is used for both negative words and the relatively disliked faces and the right key for both positive words and the relatively liked faces. Certainly your decisions should be faster than in the final part of the experiment, in which you have to press the left key for negative words and White faces and the right key for positive words and Black faces. In this case, the separate decisions are incompatible (the left key is used for negative words and relatively liked faces).

Implicit association tests such as this can be constructed to measure a wide range of biases involving age, sex, disability, attitudes to overweight or obese individuals, and so on, in addition to race. Figure 5.1 illustrates the method when testing for age bias. A substantial majority of individuals are classified as showing apparent bias by these measures. Across over 2.5 million people tested in an online version of the IAT, approximately 70 percent showed a Black-White race bias, and an even greater percentage showed an ageism bias. Black individuals showed no detectable bias, favoring neither race.¹⁷ Advocates of implicit measures like the IAT suggest that they measure very stable, long-term attitudes built up over a lifetime of social interactions embedded in our culture.

Such unconscious biases have an obvious link to the concept of microaggressions, relatively subtle acts constituting insults or slights directed toward members of disadvantaged groups. A taxi driver who drives past a Black person and picks up a White passenger is committing a microaggression, as is someone in the workplace who describes assertive females as shrill but assertive males as strong. Although less overt than direct acts of racism or sexism, the harm these cause is no less real, and many have argued that microaggressions can contribute to adverse mental health. In the past few years, a huge microaggression industry has emerged (it was the 2015 word

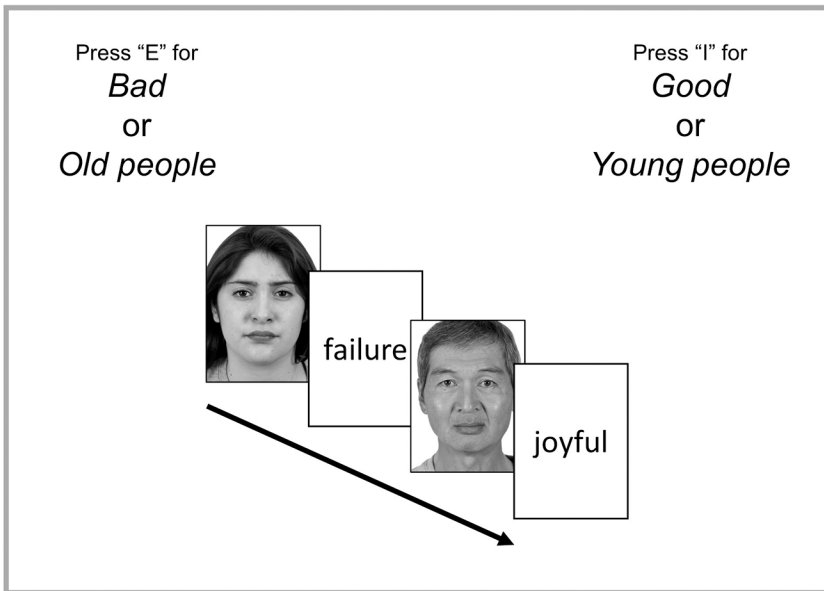


Figure 5.1

Schematic illustration of the implicit association test. This example is for measuring unconscious age bias. The faces and words illustrated here are shown one at a time on successive trials in a randomized sequence. In this consistent phase, participants press one button when a "bad" (negative) word or old face is shown, and a different button when a "good" (positive) word or young face is shown. In another (inconsistent) phase, these assignments are switched so that negative words and young faces are paired with the same button press, and the same for positive words and old faces. The difference in average speed of responding between the inconsistent and consistent phases provides the IAT measure of unconscious bias. (Faces from Debbie S. Ma, Joshua Correll, and Bernd Wittenbrink, "The Chicago Face Database: A Free Stimulus Set of Faces and Norming Data." *Behavior Research Methods* 47 (2015): 1122–1135. <https://doi.org/10.3758/s13428-014-0532-5>).

of the year, according to one source), but the status of the concept as a subject of scientific investigation is rather dubious.¹⁸ Part of the definition of microaggressions is that they lie in the eye of the beholder: if individuals feel that acts directed at them are microaggressions, then they are. There's nothing wrong with this from the perspective of using the term in our everyday cultural discourse, but a concept that cannot be objectively verified by outside observers is a poor candidate for scientific scrutiny.

Is Unconscious Bias Unconscious?

How do we know that the sorts of biases measured in the IAT or causing microaggressions are automatic and unconscious? Several justifications have been offered. Although people are often reluctant to report their overt attitudes toward disadvantaged groups, there are well-validated questionnaires for doing so. These questionnaires, which include a widely used test known as the Symbolic Racism scale, attempt to avoid such “self-presentation” or “impression management” (in social psychologists’ jargon) concerns and solicit genuine attitudes by reassuring respondents that their responses are anonymous. If the IAT were simply another method for measuring explicit attitudes, we would expect to see strong correlations when we compare individuals’ IAT scores with their explicit attitudes. This is not what is found, however. Instead, IAT scores tend to show only rather weak associations with measures of explicit attitudes, as would be expected if the IAT is measuring something distinct from conscious prejudice—that is, unconscious prejudice.¹⁹

There are, however, reasons to regard the IAT as a rather poor instrument for showing that biases can sometimes be unconscious. Indeed, two of the staunchest critics of implicit prejudice, Gregory Mitchell and Philip Tetlock, pulled no punches when they recently wrote:²⁰

It is also difficult to find a psychological construct that is so popular yet so misunderstood and lacking in theoretical and practical payoff. Scholarly discussions of prejudice fail to agree on how implicit prejudice connects to other forms of prejudice; it is unclear whether different measures of implicit prejudice measure the same thing; the meaning of “implicit” in the phrase “implicit prejudice” is contested; and implicit measures of prejudice are no better at predicting behavior, even “microaggression” (small, barely visible slights), than are traditional explicit measures of prejudice.

What motivates this coruscating criticism? There are many reasons, usually overlooked or even ignored in popular discussions of unconscious bias (in chapter 9, we discuss evidence that this biased evaluation even permeates the discussion of the IAT in introductory psychology textbooks). People seem, for instance, to have sufficient insight into their IAT responses to be able to predict them ahead of time, a capacity that surely implies a considerable degree of conscious access. When shown race IAT items—Black and White faces and positive and negative words—and asked to estimate how

easy or difficult they think sorting Black faces with positive words or sorting White faces with positive words would be, people are able to respond with considerable insight.²¹ They can also, at least to some degree, deliberately fake their responding in an IAT.²² Germans, who according to the IAT typically show prejudice against members of the minority Turkish population in Germany, can readily suppress their IAT scores and actually fake pro-Turkish attitudes when motivated to do so. The idea that the IAT is an instrument for measuring those aspects of our biases and prejudices that are beyond our conscious control is seemingly quite incorrect.

This is not the end of the difficulties faced by those who wish to use the IAT as a measure of unconscious bias, however. It appears that the test is—with a striking exception noted shortly—excessively liberal in classifying individuals as biased. Suppose we have some fairly unequivocal indicator of racial bias. Perhaps we film individuals interacting with a Black or White person and score these interactions in terms of friendliness, abruptness, physical proximity, and so on. Individuals who manifestly behave indistinguishably in interactions with a Black or White person show no tendency toward racial discrimination. And suppose we now administer the IAT to such individuals. Naturally, if it's a fair measure of bias (whether unconscious or not) it should reveal no hint of bias. But studies find something very different: that the IAT yields positive bias scores (that is, a preference for White over Black people) even in individuals who show no detectable bias in their real human interactions.²³ It's hard to interpret this as anything other than evidence of bias in a test designed to measure bias.

The exception to this liberal classification of people as biased is that the IAT reliably fails to classify men as being sexist.²⁴ In IATs designed to measure prejudice by men against women, minimal bias emerges. This is a striking and thus far unexplained anomaly. For defenders of the IAT, it must mean that there is little covert sexism in modern society, a conclusion that few would endorse. Other research shows that in contrast to explicit attitude measures, the IAT is a quite unreliable test in the sense that the scores a given individual achieves when tested twice—across even quite short intervals such as two weeks—tend not to be very consistent.²⁵

Although the IAT is by far the most widely used and evaluated test of implicit attitudes, it is far from the only one. Social psychologists have devised numerous other measures with names such as the evaluative priming and affect misattribution tests. If these tests all measure a common

underlying propensity—unconscious bias—then they should all yield similar results when large samples of people are tested on all of them. Someone who scores highly on the IAT for racism or ageism should also score highly on the other tests. But the picture is in reality rather different from this ideal. An individual classified as highly biased by the IAT is unlikely to achieve a similarly high bias score on many of the other tests. Overall, correlations between scores on these different measures tend to be quite modest.²⁶ Without compelling reason to regard any one of the tests as better than the others, this means that in reality, we have no defensible way of accurately measuring unconscious bias. Indeed the pattern of low overlap between the different tests is exactly what one might expect if in reality there's no such construct as unconscious bias and if instead each test is measuring something highly idiosyncratic.

What might such idiosyncrasies of measurement be? In the case of the IAT, there are many possibilities. It might at least in part be measuring the familiarity or salience of the target categories. For a White participant taking a race IAT, Black faces are likely to be less familiar or more salient than White ones. Perhaps it is this difference that causes faster responses for compatible than incompatible trials? It might be generally easier for people to pair familiar and good objects on the one hand and unfamiliar and bad ones on the other, rather than vice versa. This would yield the typical race IAT bias effect but with no reason to attribute it to negative reactions to other-race people. Researchers have obtained clear evidence supporting this possibility.²⁷

Significant though all these measurement problems are, an even greater concern is that implicit bias as measured by the IAT seems to have only the most tenuous connection to actual prejudicial or discriminatory behavior.²⁸ Even if the IAT were a valid and reliable measure of unconscious bias, which it plainly isn't, if this bias appears to have very little to do with overt behavior, then it's hard to argue that it's any sort of cause of how we act toward others. We've already seen that the IAT tends to classify individuals as biased even when they're behaviorally neutral. As Hart Blanton, James Jaccard, and their colleagues have argued at length and in detail, research shows that IAT scores don't do a very good job at all of predicting behavior over and above explicit reports of attitudes.²⁹

For example, if one measures the racial attitudes of a large number of individuals using standard and well-validated questionnaires such as the Symbolic Racism scale, these predict discriminatory behavior reasonably

well. If implicit bias is distinct from explicit bias, then it should follow that including scores from the IAT ought to improve even further the ability to predict discriminatory behavior, because deep-rooted unconscious biases that will fail to be detected on a self-report questionnaire will nonetheless leak through into actual behavior. But there is no research that convincingly demonstrates this. On the contrary, the added value of IAT scores is close to zero in many analyses.

The same general conclusion follows from research on the effects of interventions designed to change implicit bias. People can be given strong arguments explaining why racial bias is unacceptable, and these arguments can have some effect on reducing automatic bias, as measured by the IAT. Worryingly for defenders of the IAT, however, this research suggests that these interventions have a negligible effect on actual prejudicial behavior. Patrick Forscher and his colleagues, who conducted an extensive meta-analysis of research on this issue, noted that “there is little evidence in our data that is consistent with a causal relationship between automatically retrieved associations and behavior.”³⁰

The point is not to dismiss all of this research out of hand, and certainly not to dispute the real harm that discrimination causes to many groups. Rather, it is to emphasize that almost all of the scientific evidence for unconscious bias is controversial and open to alternative interpretations. Only when the facts have been established more convincingly one way or the other will we be able to draw firm conclusions about the specific significance of unconscious bias, in contrast to more general problems of discrimination, in the workplace and other settings.

Nor does it help when workplace training programs to counter unconscious bias lack scientific rigor. At one prestigious university (which shall remain nameless), the mandatory unconscious bias training video claims that White interviewers sit farther away from and end interviews sooner with Black candidates than White candidates, and it cites two studies documenting this claim.³¹ And the world’s oldest academic body, the Royal Society, states in its unconscious-bias video that people pay more attention to male than female voices.³²

How convincing is this evidence? Both of the studies about interviews were conducted over forty years ago and used student participants in simulated interview settings, so their relevance to modern real job interviews is

negligible. Moreover, in the study that reported that the distance between an interviewer and interviewee was greater when they were from the same race than when one was White and the other Black, it was the interviewee (not the interviewer) who chose the seating distance, so the effect cannot be evidence of interviewer bias, and the interview was about political attitudes rather than a (mock) job interview. Likewise, there is little support for the claim that people pay more attention to male than female voices. What the cited research actually found was something entirely different: that people prefer to vote for both men and women with deeper voices.³³ There is a strong impression of confirmation bias in these examples: the creators of these training videos are so convinced that unconscious bias exists that they give unjustified weight to any evidence, however dubious, that can be claimed to demonstrate it.

The sense in which unconscious bias has been assessed in this chapter—irrational preferences or behaviors—can be contrasted with a different and less controversial use of the term. This relates to ways in which we can be unaware or oblivious of the effects of our behaviors on others. To give an example, gendered language (using terms like *businessman* and *chairman*) is ubiquitous. Simply by habitually following social conventions or norms in language and culture, each of us may unwittingly reinforce disparities across gender, race, and so on. When an advertising agency uses photographs of thin models, it is contributing to our culture's idealization of a certain body shape. Many people would agree that in an equal society, such biases should be pointed out and discouraged, but there is no necessary link between using the word *chairman* and having unconscious sexist attitudes. One can follow prevalent conventions without being psychologically biased.

The societal implications of the current climate around unconscious prejudice can hardly be exaggerated. Consider the following conclusion from a recent review:

A growing body of research suggests that similar to the general US population, most HCPs [health care providers, that is, doctors] across multiple levels of training and disciplines have implicit biases against Black, Hispanic, American-Indian and dark-skinned individuals (Maina et al., 2017).

At face value, this is a gross indictment of almost the entire medical profession, stated with virtually no acknowledgment of the many reasons to be cautious, if not downright skeptical, about the validity of tests like the IAT

to accurately probe our unconscious attitudes. We can all agree that highlighting situations in which bias may occur is a worthy cause, that racism, sexism, ageism, and so on are immoral and deplorable, and that society and institutions have a duty to put policies in place to prevent discrimination and intolerance. But the evidence available to date doesn't even come close to proving that most of us walk around with unacknowledged and unconscious biases in our heads.

© 2023 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-NC-ND license.
Subject to such license, all rights are reserved.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Newell, Benjamin R., 1972– author. | Shanks, David R.

Title: Open minded : searching for truth about the unconscious mind /
Ben R. Newell and David R. Shanks.

Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes
bibliographical references and index.

Identifiers: LCCN 2022038725 (print) | LCCN 2022038726 (ebook) |
ISBN 9780262546195 (paperback) | ISBN 9780262375368 (epub) |
ISBN 9780262375375 (pdf)

Subjects: LCSH: Subconsciousness. | Thought and thinking. | Self-consciousness
(Awareness)

Classification: LCC BF315 .N479 2023 (print) | LCC BF315 (ebook) |
DDC 154.2—dc23/eng/20230316

LC record available at <https://lcn.loc.gov/2022038725>

LC ebook record available at <https://lcn.loc.gov/2022038726>