


CHAPTER 4: CONCEPTUAL ILLUSIONS

In this chapter we will consider some fallacies and peculiarities of the way we reason. We'll briefly consider the fallacy of denying the evidence and then confirmation bias, the paradox of choice, the conjunction fallacy and implicit bias.

1 DENYING THE EVIDENCE



2 CONFIRMATION BIAS

"It is the peculiar and perpetual error of the human understanding to be more moved and excited by affirmatives than by negatives."

--Francis Bacon

Confirmation bias refers to a type of selective thinking whereby one tends to notice and to look for what confirms one's beliefs, and to ignore, not look for, or undervalue the relevance of what contradicts one's beliefs. For example, if you believe that during a full moon there is an increase in admissions to the emergency room where you work, you will take notice of admissions during a full moon, but be inattentive to the moon when admissions occur during other nights of the month. A tendency to do this over time unjustifiably strengthens your belief in the relationship between the full moon and accidents and other lunar effects.

This tendency to give more attention and weight to data that support our beliefs than we do to contrary data is especially pernicious when our beliefs are little more than prejudices. If our beliefs are firmly established on solid evidence and valid confirmatory experiments, the tendency to give more attention and weight to data that fit with our beliefs should not lead us astray as a rule. Of course, if we become blinded to evidence truly refuting a favored hypothesis, we have crossed the line from reasonableness to closed-mindedness.


Numerous studies have demonstrated that people generally give an excessive amount of value to confirmatory information, that is, to positive or supportive data. The "most likely reason for the excessive influence of confirmatory information is that it is easier to deal with cognitively" (Gilovich 1993). It is much easier to see how a piece of data supports a position than it is to see how it might count against the position. Consider a typical ESP experiment or a seemingly clairvoyant dream: Successes are often unambiguous or data are easily massaged to count as successes, while negative instances require intellectual effort to even see them as negative or to consider them as significant. The tendency to give more attention and weight to the positive and the confirmatory has been shown to influence memory. When digging into our memories for data relevant to a position, we are more likely to recall data that confirms the position (ibid.).

Researchers are sometimes guilty of confirmation bias by setting up experiments or framing their data in ways that will tend to confirm their hypotheses. They compound the problem by proceeding in ways that avoid dealing with data that would contradict their hypotheses. For example, some parapsychologists used to engage in optional starting and stopping in their ESP research. Experimenters might avoid or reduce confirmation bias by collaborating in experimental design with colleagues who hold contrary hypotheses, as Richard Wiseman (skeptic) and Marilyn Schlitz (proponent) have done.* Individuals have to constantly remind themselves of this tendency and actively seek out data contrary to their beliefs. Since this is unnatural, it appears that the ordinary person is doomed to bias.

3 HOW TO FALSIFY A CONDITIONAL

Do you have Confirmation Bias? The results of the following psychological experiment suggest that most of us do insofar as when asked to test a hypothesis we look for evidence that confirms it and ignore data that would falsify it.

Consider the following hypothesis: If a card has a vowel on one side then it has an even number on the other. You may turn over cards to see what is on the other side. We assume that each card has a letter on one side and a number on the other. In the illustration, the first two cards have their letter sides facing up and an unknown number on their other sides. The last two have their number sides up with the letters, which we can't yet see, facing downwards.



How do you determine the most effective means for deciding whether the hypothesis is correct or incorrect? Which cards would you turn—assuming that you want to turn no more than you need to establish whether the hypothesis is true or false?

Most subjects thought that only the A and 4 cards needed to be turned. If the A card had an even number on its other side that would support (though not conclusively confirm) the hypothesis. Similarly, if the 4 card showed a vowel on the other side that would be consistent with the hypothesis. Of course, if it had a consonant on the other side that wouldn't show anything. Remember, this is an hypothesis about vowel cards. It doesn't say anything about consonant cards so if there were a consonant on the other side of the 4 card it wouldn't tell you anything.

Subjects were at least dimly aware of this because most realized that turning the D card would be a waste. The hypothesis, to repeat, says nothing about what we should find on consonant, i.e. non-vowel cards. So, whatever turned up on the other side of the D card just wouldn't matter.

Subjects however failed to realize that to test the hypothesis it was crucial to turn over the 7 card because if there were a vowel on the other side that would show conclusively that the hypothesis was false. Suppose there were an I on the flip side of the 7 card. The hypothesis says:

If a card has a vowel on one side then it has an even number on the other.

In the circumstances described the 7-card, a.k.a. the I-card, would have a vowel on one side but would not have an even number on the other—contrary to the hypothesis!

A hypothesis like this is an if-then statement—what we call a *conditional* in logic. We call the if clause the *antecedent* and the then clause the *consequent*. So the antecedent of this conditional is: “a card has a vowel on one side”; the consequent is: “it has an even number on the other [side]” To show conclusively that a conditional is false, we come up with a case in which the antecedent is true and the consequent is false. Such a case is a counterexample to the conditional. So, the case of a card with an I on one side and a 7 on the other is a counterexample to our hypothesis.

4 CONJUNCTION FALLACY

Which of the following events is most likely to occur within the next year:

The United States will withdraw all troops from Iraq.

The United States will withdraw all troops from Iraq and bomb Iranian nuclear facilities.

Answer: <http://www.fallacyfiles.org/conjunct.html> - Answer

The probability of a conjunction is never greater than the probability of its conjuncts. In other words, the probability of two things being true can never be greater than the probability of one of them being true, since in order for both to be true, each must be true. However, when people are asked to compare the probabilities of a conjunction and one of its conjuncts, they sometimes judge that the conjunction is more likely than one of its conjuncts. This seems to happen when the conjunction suggests a scenario that is more easily imagined than the conjunct alone.

Interestingly, psychologists Kahneman and Tversky discovered in their experiments that statistical sophistication made little difference in the rates at which people committed the conjunction fallacy.

Interestingly, psychologists Kahneman and Tversky discovered in their experiments that statistical sophistication made little difference in the rates at which people committed the conjunction fallacy. One of Kahneman and Tversky's tests is known as "the Linda Problem"¹:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Is it more likely that Linda is a bank teller, or a bank teller and feminist? Of course, it is more likely that she is the conjunct than the conjunction. However, the description of Linda given in the problem fits the stereotype of a feminist, whereas it doesn't fit the stereotypical bank teller. Because it is easy to imagine Linda as a feminist, people may misjudge that she is more likely to be both a bank teller and a feminist than a bank teller.



¹ Amos Tversky & Daniel Kahneman, "Judgments of and by Representativeness", in *Judgment Under Uncertainty: Heuristics and Biases*, Kahneman, Paul Slovic, and Tversky, editors (1985), pp. 84-98.

4.1 MORAL: IN WEAKNESS THERE IS STRENGTH!

An and-statement is called a *conjunction*. Its parts are called *conjuncts*. The statement that Linda is a bank teller and is a feminist activist is stronger than either of its conjuncts in the sense that it tells us more about the world—it commits us to more if we believe it and it entails more. Strength however is a mixed blessing: the more we say, the more likely we are to go wrong! In this respect, there is strength in weakness—the less we commit to the less likely we are to go wrong! The more careful we are about what we claim, the less we commit to, the more defensible our position will be!

5 THE PARADOX OF CHOICE: CAN WE HAVE TOO MUCH OF A GOOD THING?



Prima facie, the Good Life is simply a matter of being able to get what we want—whatever it is. Intuitively, the more real options we have the better off we are since the more options we have the more likely we are to satisfy our preferences.

Empirical research on “happiness” however suggests that we are generally rather poor judges of what states of affairs will make us happiest. Citing the results of this research and work on “bounded rationality” by Herbert Simon and others, popularizers, like Barry Schwartz, note that there is a “paradox of choice” insofar as the availability of a wide range of options does not typically make us better off and quite often makes us worse off. We are bewildered and confused by a wide range alternatives, agonize over making the best possible choices, invest too heavily in searching for the optimal solution, and end up worse off than we would be if, facing a narrower range of options, we had merely aimed to satisfice.

All other things being equal the more choices we have the better off we are. But all other things are rarely equal...

5.1 MORE IS NOT ALWAYS BETTER

Does having more options always make us better off?

According to the growing empirical literature on “happiness” individuals who live in relatively restricted circumstances may be as happy, or happier, than those with much wider scope for choice. Wealth is the permanent possibility of preference satisfaction—the capability of satisfying a wide range of desires. Empirical research however suggests that while happiness increases with increasing wealth to a point, after that point more money does not make much difference. Individuals in some developing countries who are, by American standards, poor are

as happy as Americans who, with vast purchasing power, have a much wider range of options. Moreover privileged individuals whose options are curtailed quickly adjust. People who predict that they would be thoroughly miserable if they were incapacitated by illness or injury bounce back. Within months after losing limbs, for example, many amputees report levels of happiness comparable to those they enjoyed prior to amputation.

5.2 TOO MUCH OF A GOOD THING

Can having more options make us worse off?

Moreover empirical data suggest also that the availability of a great many options, particularly those we regard as desirable, can make us worse off. Drawing on the work of Daniel Kahneman and Amos Tversky in behavioral economics and the results of empirical research by psychologists, market researchers and decision scientists, Barry Schwartz argues that "we would be better off if we embraced certain voluntary constraints on our freedom of choice, instead of rebelling against them" and promotes other doctrines which "fly in the face of conventional wisdom that the more choices people have, the better off they are."² According to Schwartz, more choices may in a range of cases make us worse off.

In the cases that concern Schwartz there are typically a plethora of options that confuse, befuddle and, occasionally, incapacitate consumers facing a bewildering array of products. Paradoxically, Schwartz notes, the more choices consumers have the less likely they are to be satisfied. Furthermore, becoming informed and coolly deliberating only makes things worse as rational choosers, aiming to maximize preference satisfaction waste time, effort and resources on research and deliberation.

5.3 WHEN CHOICE IS DEMOTIVATING

One study was set in a gourmet food store in an upscale community where, on weekends, the owners commonly set up sample tables of new items. When researchers set up a display featuring a line of exotic, high-quality jams, customers who came by could taste samples, and they were given a coupon for a dollar off if they bought a jar. In one condition of the study, 6 varieties of the jam were available for tasting. In another, 24 varieties were available. In either case, the entire set of 24 varieties were available for purchase. The large array of jams attracted more people to the table than the small array, though in both cases people tasted about the same number of jams on the average. When it came to buying, however, a huge difference became evident. Thirty percent of the people exposed to the small array of jams actually bought a jar; only 3 percent of those exposed to the large array of jams did so.³

² Barry Schwartz. *The Paradox of Choice*. New York: HarperCollins Publishers inc., 2004), p. 5.

³ Schwartz 2004, pp. 19-20.



In cases like this however it is not the availability of a wide range of desirable options that confuses consumers but their salience. We are confused by a proliferation of visual “noise” — and noise can be dealt with. Astute shopkeepers display selected samples of representative products even where a much wider range is available: the owners of the gourmet food shop normally restricted the display at which samples were offered to half a dozen jams even though consumers in fact had a wider choice. In cases like this it is not the availability of many options but their intrusive visibility that befuddles us.

When shopkeepers don’t filter out noise we can do it for ourselves. Schwartz notes that even faced with the same range of options, consumers adopt very different strategies when it comes to deciding which to pursue.

5.4 MAXIMIZERS AND SATISFICERS

Maximizers need to be assured that every purchase or decision was the best that could be made...The alternative to maximizing is to be a satisficer. To satisfice is to settle for something that is good enough and not worry about the possibility that there might be something better...In the end...[maximizers] are likely to get less satisfaction out of the exquisite choices they make than will satisficers...When Nobel Prize-winning economist and psychologist Herbert Simon initially introduced the idea of 'satisficing' in the 1950s, he suggested that when all the costs (in time, money, and anguish) involved in getting information about all the options are factored in, satisficing is in fact the maximizing strategy. In other words, the best people can do, all things considered, is to satisfice.⁴

Since I know that satisficing is more conducive to preference satisfaction than attempting to maximize I can, rationally, make the decision to satisfice. If this is correct, then Maximizers, who insist on giving full consideration to all available options, are irrational—they have adopted a bad strategy for securing preference satisfaction and, more often than not, fail. Their problem

⁴ Schwartz, pp. 78-79. Schwartz includes in his book a psychological survey that readers can take to determine where they fall on the maximizing-satisficing continuum. After giving the survey to "several thousand people" he notes that the high score was 75 and the low was 25.

is not the range of options available to them but the way in which they deal with these options. Satisficers faced with the same alternatives do splendidly: they adopt stopping rules and, in circumstances where the range of options is confusing, filter out noise by means of self-imposed constraints.⁵ Maximizers, whether out of ignorance or through weakness of will shop until they drop and usually end up worse off even when they manage to get what they regard as the best possible products.

To understand why, it is crucial to recognize what the objects of consumer choice actually are. Arguably they are not discrete, spatio-temporally bounded consumer products—canned goods, cars or boxes of laundry detergent—but states of being and doing which include such intangibles as leisure and comfort. When a Maximizer sorts through every trinket on the table or hits every website looking for the best possible item at the lowest possible price, comparison shops, reads labels and checks out consumer magazines she is choosing a state of being and doing that includes not only the possession of the product she eventually buys but the trips to stores all around town, the time and effort spent reading consumer magazines, and the process of deliberation. The price she pays also includes substantial opportunity costs. Even if she succeeds in choosing the product that she would prefer all other things being equal, all other things are not equal and she might have done better overall to have grabbed a merely satisfactory product early on without research or comparison shopping, and spent her time doing more interesting things.

Having a wide range of options per se can never make us worse off. But believing that we do can cause trouble—and for Maximizers quite often does. The fault however is not in our stars but in ourselves: maximizing strategies are often doomed to failure whereas satisficing strategies in most cases are the most effective way to maximize!

6 IMPLICIT BIAS

6.1 THE IAT TEST

AT 4 O'CLOCK ON A RECENT WEDNESDAY AFTERNOON, a 34-year-old white woman sat down in her Washington office to take a psychological test. Her office decor attested to her passion for civil rights -- as a senior activist at a national gay rights organization, and as a lesbian herself, fighting bias and discrimination is what gets her out of bed every morning. A rainbow flag rested in a mug on her desk.

⁵ See, e.g. David Schmitz, "Satisficing as a Humanly Rational Strategy" in Byron, *Satisficing and Maximizing*, pp. 33 ff. Satisficers, recognizing that their pursuit of what Schmitz calls "local goals" may compromise their ability to attain "global goals," compartmentalize. They impose budget constraints on themselves to limit their choice set when it comes to the pursuit of local goals, and seek local optima within those constraints.

The woman brought up a test on her computer from a Harvard University Web site. It was really very simple: All it asked her to do was distinguish between a series of black and white faces. When she saw a black face she was to hit a key on the left, when she saw a white face she was to hit a key on the right. Next, she was asked to distinguish between a series of positive and negative words. Words such as "glorious" and "wonderful" required a left key, words such as "nasty" and "awful" required a right key. The test remained simple when two categories were combined: The activist hit the left key if she saw either a white face or a positive word, and hit the right key if she saw either a black face or a negative word.

Then the groupings were reversed. The woman's index fingers hovered over her keyboard. The test now required her to group black faces with positive words, and white faces with negative words. She leaned forward intently. She made no mistakes, but it took her longer to correctly sort the words and images. Her result appeared on the screen, and the activist became very silent. The test found she had a bias for whites over blacks.

"It surprises me I have any preferences at all," she said. "By the work I do, by my education, my background. I'm progressive, and I think I have no bias. Being a minority myself, I don't feel I should or would have biases."

Although the activist had initially agreed to be identified, she and a male colleague who volunteered to take the tests requested anonymity after seeing their results. The man, who also is gay, did not show a race bias. But a second test found that both activists held biases against homosexuals -- they more quickly associated words such as "humiliate" and "painful" with gays and words such as "beautiful" and "glorious" with heterosexuals. If anything, both activists reasoned, they ought to have shown a bias in favor of gay people. The man's social life, his professional circle and his work revolve around gay culture. His home, he said, is in Washington's "gayborhood."

"I'm surprised," the woman said. She bit her lip. "And disappointed."

"THE IMPLICIT ASSOCIATION TEST measures the thumbprint of the culture on our minds," says Banaji, one of three researchers who developed the test and its most ardent proponent. "If Europeans had been carted to Africa as slaves, blacks would have the same beliefs about whites that whites now have about blacks."

The results of the millions of tests that have been taken anonymously on the Harvard Web site and other sites hint at the potential impact of the research. Analyses of tens of thousands of tests found 88 percent of white people had a pro-white or anti-black implicit bias; nearly 83 percent of heterosexuals showed implicit biases for straight people over gays and lesbians; and more than two-thirds of non-Arab, non-Muslim volunteers displayed implicit biases against Arab Muslims. Overall, according to the researchers, large majorities showed biases for Christians over Jews, the rich over the poor, and men's careers over women's careers. The results contrasted sharply with what most people said about themselves -- that they had no biases. The

tests also revealed another unsettling truth: Minorities internalized the same biases as majority groups. Some 48 percent of blacks showed a pro-white or anti-black bias; 36 percent of Arab Muslims showed an anti-Muslim bias; and 38 percent of gays and lesbians showed a bias for straight people over homosexuals.

As the tests have been refined, replicated and reinterpreted over the past decade, they have challenged many popular notions -- beginning with the increasingly common assertion that discrimination is a thing of the past. The research has also upset notions of how prejudice can best be addressed. Through much of the 20th century, activists believed that biases were merely errors of conscious thought that could be corrected through education. This hopeful idea is behind the popularity of diversity training. But Banaji suggests such training relies on the wrong idea of how people form biases.

There is likely a biological reason people so quickly make assumptions -- good or bad -- about others, Banaji says. The implicit system is likely a part of the "primitive" brain, designed to be reactive rather than reasoned. It specializes in quick generalizations, not subtle distinctions. Such mental shortcuts probably helped our ancestors survive. It was more important when they encountered a snake in the jungle to leap back swiftly than to deduce whether the snake belonged to a poisonous species. The same mental shortcuts in the urban jungles of the 21st century are what cause people to form unwelcome stereotypes about other people, Banaji says. People revert to the shortcuts simply because they require less effort. But powerful as such assumptions are, they are far from permanent, she says. The latest research, in fact, suggests these attitudes are highly malleable. Such reassurance has not assuaged test takers, who are frequently shocked by their results.

In another study by psychologist Robert W. Livingston at the University of Wisconsin, Poehlman says, volunteers were given details of a crime in which a Milwaukee woman had been assaulted, suffered a concussion and required several stitches. In this case, Poehlman says, some volunteers were told the perpetrator had been proven to be David Edmonds from Canada. Others were told the guilty perpetrator was Juan Luis Martinez from Mexico. Volunteers were asked what length of sentence was appropriate for the crime: Bias scores against Hispanics on the implicit tests tended to predict a longer sentence for the Mexican. An implicit attitude "doesn't control our behavior in a be-all and end-all kind of way, but it flavors our behavior in a pretty consistent way," says Poehlman.

In perhaps the most dramatic real-world correlate of the bias tests, economists at the Massachusetts Institute of Technology and the University of Chicago recently sent out 5,000 résumés to 1,250 employers who had help-wanted ads in Chicago and Boston. The résumés were culled from Internet Web sites and mailed out with one crucial change: Some applicants were given stereotypically white-sounding names such as Greg; others were given black-sounding names such as Tyrone. Interviews beforehand with human resources managers at many companies in Boston and Chicago had led the economists to believe that black applicants

would be more likely to get interview calls: Employers said they were hungry for qualified minorities and were aggressively seeking diversity. Every employer got four résumés: an average white applicant, an average black applicant, a highly skilled white applicant and a highly skilled black applicant.

The economists measured only one outcome: Which résumés triggered callbacks? To the economists' surprise, the résumés with white-sounding names triggered 50 percent more callbacks than résumés with black-sounding names. Furthermore, the researchers found that the high-quality black résumés drew no more calls than the average black résumés. Highly skilled candidates with white names got more calls than average white candidates, but lower-skilled candidates with white names got many more callbacks than even highly skilled black applicants.



"Fifty percent? That's huge," says Sendhil Mullainathan, an economist who led the study and who recently moved to Harvard to work with Banaji. Human resources managers were stunned by the results, he says. Explicit bias, says Mullainathan, can occur not only without the intent to discriminate, but despite explicit desires to recruit minorities. Implicit attitudes need only sway a few decisions to have large impact, he says. For example, if implicit bias caused a recruiter to set one résumé aside, it could be just one of 100 decisions the recruiter made that day. Collectively, however, such decisions can have dramatically large consequences.

SAJ-NICOLE JONI WAS THE FIRST WOMAN TO BE HIRED AS AN APPLIED MATHEMATICS PROFESSOR AT MIT. It was 1977, and there were no women's bathrooms in her building. Joni was not particularly surprised. She had battled obstacles all her life. When she first declared -- at age 12 -- that she was going to be a mathematician, her announcement evoked gales of laughter at a family gathering. But opposition only made her more determined. After a successful stint at MIT, Joni worked for Microsoft and then launched a successful business consulting firm called the Cambridge International Group Ltd. Her recent book, *The Third Opinion*, stresses the importance of seeking diverse points of view.



Joni was recently introduced to Banaji and expressed interest in taking the Implicit Association Test. Like most volunteers, she did not think she had biases and believed strongly in "meeting people as they are, without looking at the color of their skin." Given Joni's background, Banaji thought it would be interesting for her to take a bias test that examined whether Joni associated men or women with careers in science. Most people find it easier to associate men

with the sciences -- but Joni was clearly not most people. The test came up on the screen. Joni's fingers, trained for many years on the piano, flew as she classified a number of words such as "husband," "father," "mother" and "wife" between "male" and "female" groups. She then grouped words such as "chemistry," "history," "astronomy" and "music" under "science" or "liberal arts." The computer then asked her to group "male" with "science" and "female" with "liberal arts."

When the groupings were reversed, Joni had to group "male" words with "liberal arts," and "female" words with various disciplines in science. She made a mistake in classifying "uncle." She hesitated over "astronomy" and made a second mistake in classifying "physics." The results popped up: "Your data show a strong association between science and Male relative to Female." Joni's fingers tapped the table in frustration. "I fought for women to be scientists all my life," she said, incredulous. Banaji nodded sympathetically. Her own results on this test were similar.

The IAT research shows that hostility is not needed for discrimination to occur. Women and minorities can just as easily harbor biases, absorbed from the larger culture, that can lead them to discriminate against people like themselves...[H]uman beings filter what they see through the lenses of their own expectations. People believe they are acting rationally, but numerous psychological tests prove that subtle cues influence people all the time without their knowledge.

Banaji believes that conscious efforts are needed to fight what she calls ordinary prejudice, the primitive brain filtering the world through its biased lenses without the conscious part of the brain being aware of it. Since Mullainathan found startling differences in his résumé study, he says, he has come to believe that personal identifiers should be removed from résumés to make evaluations more fair. Another area highly prone to implicit biases is job interviews, says Max Bazerman of Harvard Business School. "What you need to do is look at objective measures separate from the interview."

Banaji and Kang believe the IAT can be used as one measure to determine when affirmative action policies ought to be ended. Rather than pick an arbitrary amount of time -- Supreme Court Justice Sandra Day O'Connor recently suggested 25 years -- the researchers asked whether such policies should expire when implicit tests show that people are really evaluating others without bias. Banaji and Kang are less interested in using affirmative action to redress historical wrongs -- they argue it is essential to fight discrimination still taking place today.

"Mind bugs operate without us being conscious of them," Banaji says. "They are not special things that happen in our heart because we are evil."

6.2 BEYOND THE IAT

The IAT seems a little gimmicky, but the results are corroborated in empirical studies by social scientists. In the following January 2015 article, Harvard economist Sendhil Mullainathan summarizes the results.

6.3 RACIAL BIAS, EVEN WHEN WE HAVE GOOD INTENTIONS

By Sendhil Mullainathan

Arguments about race are often heated and anecdotal. As a social scientist, I naturally turn to empirical research for answers. As it turns out, an impressive body of research spanning decades addresses just these issues — and leads to some uncomfortable conclusions and makes us look at this debate from a different angle.

The central challenge of such research is isolating the effect of race from other factors. For example, we know African-Americans earn less income, on average, than whites. Maybe that is evidence that employers discriminate against them. But maybe not. We also know African-Americans tend to be stuck in neighborhoods with worse schools, and perhaps that — and not race directly — explains the wage gap. If so, perhaps policy should focus on place rather than race, as some [argue](#).

But we can isolate the effect of race to some degree. A [study](#) I conducted in 2003 with [Marianne Bertrand](#), an economist at the University of Chicago, illustrates how.⁶ We mailed thousands of résumés to employers with job openings and measured which ones were selected for callbacks for interviews. But before sending them, we randomly used stereotypically African-American names (such as “Jamal”) on some and stereotypically white names (like “Brendan”) on others.

The same résumé was roughly 50 percent more likely to result in callback for an interview if it had a “white” name. Because the résumés were statistically identical, any differences in outcomes could be attributed only to the factor we manipulated: the names.

Other studies have also examined race and employment. In a 2009 study, Devah Pager, Bruce Western and Bart Bonikowski, all now sociologists at Harvard, [sent](#) actual people to apply for low-wage jobs. They were given identical résumés and similar interview training. Their sobering finding was that African-American applicants with no criminal record were offered jobs at a rate as low as white applicants who had criminal records.

⁶ USD students can access the complete study here: <http://0-web.b.ebscohost.com.sally.sandiego.edu/ehost/pdfviewer/pdfviewer?vid=2&sid=3af1cc6e-1a16-4ca0-8551-ab69f6b4c51b%40sessionmgr112&hid=125>

These kinds of methods have been used in a variety of research, especially in the last 20 years. Here are just some of the general findings:

- When doctors [were shown](#) patient histories and asked to make judgments about heart disease, they were much less likely to recommend cardiac catheterization (a helpful procedure) to black patients — even when their medical files were statistically identical to those of white patients.
- When whites and blacks were [sent](#) to bargain for a used car, blacks were offered initial prices roughly \$700 higher, and they received far smaller concessions.
- Several studies found that sending emails with stereotypically black names in response to apartment-rental ads on Craigslist elicited fewer responses than sending ones with white names. A regularly repeated [study](#) by the federal Department of Housing and Urban Development sent African-Americans and whites to look at apartments and found that African-Americans were shown fewer apartments to rent and houses for sale.
- White state legislators were found to be less [likely](#) to respond to constituents with African-American names. This was true of legislators in both political parties.
- [Emails](#) sent to faculty members at universities, asking to talk about research opportunities, were more likely to get a reply if a stereotypically white name was used.
- Even eBay [auctions](#) were not immune. When iPods were auctioned on eBay, researchers randomly varied the skin color on the hand holding the iPod. A white hand holding the iPod received 21 percent more offers than a black hand.

The criminal justice system — the focus of current debates — is harder to examine this way. [One study](#), though, found a clever method. The pools of people from which jurors are chosen are effectively random. Analyzing this natural experiment revealed that an all-white jury was 16 percentage points more likely to convict a black defendant than a white one, but when a jury had one black member, it convicted both at the same rate.

I could go on, but hopefully the sheer breadth of these findings impresses you, as it did me.

There are some counterexamples: Data show that some places, like elite colleges, [most likely](#) do favor minority applicants. But this evidence underlies that a helping hand in one area does not preclude harmful shoves in many other areas, including ignored résumés, unhelpful faculty members and reluctant landlords.

But this widespread discrimination is not necessarily a sign of widespread *conscious* prejudice. When our own résumé study came out, many human-resources managers told us

they were stunned. They prized creating diversity in their companies, yet here was evidence that they were doing anything but. How was that possible? To use the language of the psychologist Daniel Kahneman, we think both fast and slow. When deciding what iPod to buy or which résumé to pursue, we weigh a few factors deliberately (“slow”). But for hundreds of other factors, we must rely on intuitive judgment — and we weigh these unconsciously (“fast”). Even if, in our slow thinking, we work to avoid discrimination, it can easily creep into our fast thinking. Our snap judgments rely on all the associations we have — from fictional television shows to news reports. They use stereotypes, both the accurate and the inaccurate, both those we would want to use and ones we find repulsive.

We can’t articulate why one seller’s iPod photograph looks better; dozens of factors shape this snap judgment — and we might often be distraught to realize some of them. If we could make a slower, deliberate judgment we would use some of these factors (such as the quality of the photo), but ignore others (such as the color of the hand holding the iPod). But many factors escape our consciousness.

This kind of discrimination — crisply articulated in a 1995 [article](#) by the psychologists Mahzarin Banaji of Harvard and Anthony Greenwald of the University of Washington — has been studied by dozens of researchers who have documented implicit bias outside of our awareness. The key to “fast thinking” discrimination is that we all share it. Good intentions do not guarantee immunity. One [study](#) published in 2007 asked subjects in a video-game simulation to shoot at people who were holding a gun. (Some were criminals; some were innocent bystanders.) African-Americans were shot at a higher rate, even those who were not holding guns.

Ugly pockets of conscious bigotry remain in this country, but most discrimination is more insidious. The urge to find and call out the bigot is powerful, and doing so is satisfying. But it is also a way to let ourselves off the hook. Rather than point fingers outward, we should look inward — and examine how, despite best intentions, we discriminate in ways big and small.

Sendhil Mullainathan is a professor of economics at Harvard. Follow him on Twitter at [@m_sendhil](#).

NYTimes: Economic View, Jan 3, 2015

7 APPENDIX

TABLE 1—MEAN CALLBACK RATES BY RACIAL SOUNDINGNESS OF NAMES

	Percent callback for White names	Percent callback for African-American names	Ratio	Percent difference (<i>p</i> -value)
Sample:				
All sent resumes	9.65 [2,435]	6.45 [2,435]	1.50	3.20 (0.0000)
Chicago	8.06 [1,352]	5.40 [1,352]	1.49	2.66 (0.0057)
Boston	11.63 [1,083]	7.76 [1,083]	1.50	4.05 (0.0023)
Females	9.89 [1,860]	6.63 [1,886]	1.49	3.26 (0.0003)
Females in administrative jobs	10.46 [1,358]	6.55 [1,359]	1.60	3.91 (0.0003)
Females in sales jobs	8.37 [502]	6.83 [527]	1.22	1.54 (0.3523)
Males	8.87 [575]	5.83 [549]	1.52	3.04 (0.0513)

Notes: The table reports, for the entire sample and different subsamples of sent resumes, the callback rates for applicants with a White-sounding name (column 1) an an African-American-sounding name (column 2), as well as the ratio (column 3) and difference (column 4) of these callback rates. In brackets in each cell is the number of resumes sent in that cell. Column 4 also reports the *p*-value for a test of proportion testing the null hypothesis that the callback rates are equal across racial groups.