

Fall 2017

## Police in America: Ensuring Accountability and Mitigating Racial Bias Feat. Professor Destiny Peery

---

### Recommended Citation

*Police in America: Ensuring Accountability and Mitigating Racial Bias Feat. Professor Destiny Peery*, 11 Nw. J. L. & Soc. POL'Y. 354 ().  
<http://scholarlycommons.law.northwestern.edu/njlsj/vol11/iss4/1>

This Conference Proceeding is brought to you for free and open access by Northwestern University School of Law Scholarly Commons. It has been accepted for inclusion in Northwestern Journal of Law & Social Policy by an authorized editor of Northwestern University School of Law Scholarly Commons.

**NORTHWESTERN  
JOURNAL OF LAW AND SOCIAL POLICY  
NINTH ANNUAL SYMPOSIUM**

**Police in America: Ensuring Accountability and  
Mitigating Racial Bias**

UNDERSTANDING AND OVERCOMING IMPLICIT BIAS held at Northwestern Pritzker School of Law, Thorne Auditorium, 375 East Chicago Avenue, Chicago, Illinois, on the 13<sup>th</sup> day of November, A.D. 2015.

FEATURED SPEAKER: Professor Destiny Peery; Introduction by Professor Locke Bowman.

PROFESSOR BOWMAN: Good morning and welcome, everyone. My name is Locke Bowman. I am the director of the MacArthur Justice Center and a professor here at the law school. And I'm pleased and honored to welcome you all to a day of conversation about what works; about how litigation, grass roots activism, and policy reform are making a difference in combatting our unacceptable level of police-on-citizen violence. That is the reality that motivates this conference. And that is the starting point for our discussion today.

We doubt that anyone disputes this. Indeed, what is most remarkable about the aftermath of Michael Brown's shooting a year ago last August is not the avalanche of commentary and protest in the streets. It is not the glare of media attention, nor is it the sustained salience of that attention.

What is most remarkable is that this great outpouring arose from something so senseless, and yet, so ordinary: a young African-American man killed in an encounter with law enforcement.

The government cannot tell us how frequently police kill citizens, but we know enough to know that the killing of a young black man on the streets of this country is an everyday occurrence.

According to the Guardian's "The Counted" Project, there have been nearly a thousand citizen fatalities at the hands of police in the first ten and a half months of this year.

Those most at risk are African-Americans, whose fatality rate, adjusted for their share of the population, is three times that of whites and other races. Police have been charged with crimes of any kind in less than 1% of these fatal encounters. Rates of police discipline for these homicides aren't known, but I will wager that they are even lower.

And, of course, police homicides are but the tip of an iceberg that includes tens of thousands of arrests, more violent than necessary; countless discretionary so-called investigative stops of motorists, mainly black and brown, for trivial or nonexistent infractions; the daily racism and indignities of stop and frisk; the unjustified so-called mob action; arrests of young black men and boys for the so-called crime of standing on a street corner.

Our conversation today will be about what we can do to mitigate the racial bias that contributes to this misconduct and what we can do to make police accountable when they victimize citizens.

It is fitting that we should begin our discussion today with a talk about why why African-Americans are so disproportionately the victims of police violence.

Destiny Peery frames the issue this way: What are the psychological mechanisms, including unacknowledged bias and threats produced by stereotyping, that interact with police training, the social context of policing, and existing policies and procedures to foster abuses of authority, the use of violence and sour relationships between police and their communities?

Destiny is the perfect person to tackle this question. She is a J.D. She is also a Ph.D., both from Northwestern. She holds a joint appointment in this law school and in our University's Department of Psychology. Destiny has written and published extensively about race, people's perception of race, and the effects of those perceptions.

Among many other projects, she is one of the authors of a prominent and much-cited study about the way we read mixed-race individuals. Are they biracial, or are they black?

Destiny is a popular teacher here. She is a sought-after lecturer. She is lucid and smart. And I, for one, am very much looking forward to her talk.

Please join me in welcoming Destiny Peery.

**PROFESSOR PEERY:** That was a very kind introduction.

I'm here to provide a foundation for thinking about bias in policing from a psychological standpoint. The starting point for me, as a lawyer, a law professor, and a psychologist, is: what is bias? What are the underlying mechanisms that inform not only policing, but also, more generally, how we interact with and perceive others in the world around us?

Of course, where we're starting in the context of this symposium—and it's already been referenced multiple times in this setting—is the Ferguson spotlight. We have a situation in Ferguson, with the killing of Michael Brown, that brings to nationwide attention the idea of bias in policing, including its effects on criminal justice disparities more broadly—things like mass incarceration, disparate impact of arrests, and disparate suspicion trained on certain types of individuals over others. What can we do about this? Is this about implicit bias? Is this about explicit bias? Do we address it through training? Do we address it through some other means? And that's part of the foundation for this type of symposium, but also a foundation for what is now a national discourse on bias in policing.

I want to highlight that none of these issues are new. Questions about implicit bias, explicit bias, the impact of bias on how we perceive and react to other people: they're not new. They didn't first appear when we saw Michael Brown lying in the street that hot August afternoon. The idea of implicit bias in psychology has been around for decades. And certainly explicit bias, or prejudice, and discrimination have been around for probably as long as humans have been around. But current events give us a specific context in which to focus on the role of bias in our interactions in society, and in particular, given our setting here in a law school, the effect of bias on the legal system.

Underlying our understanding of current events—and this is why I'm starting the symposium off here—is the psychology of bias. Once we abstract from a particular context,

we can focus on what's going on in our brains—what's going on in our minds—when we stereotype another person, think another person is *less than*, or treat another person in a way that's dehumanizing or that in some way harms them.

I usually start off this introduction with a discussion of the triad of stereotyping, prejudice, and discrimination, because we tend to hear those three things talked about together. From a psychological standpoint, they are separate but related concepts that can have different consequences for our interactions with other people.

We tend to hear the most, day to day, about stereotypes, and stereotyping is in some ways the most concrete concept of the three. Stereotypes are beliefs about certain groups or types of people having certain attributes. For example, thinking that black people are good at sports, Asian people are good at math, women are caring, men are aggressive; those are all stereotypes. They're ideas that we attach to groups. A stereotype isn't about whether any one man, woman, black person, Asian person, et cetera has this this or that trait. Instead, a stereotype takes it as given that by virtue of being a member of a group, you have a certain characteristic or trait.

Stereotypes can be both descriptive, meaning they focus on what certain types of people *are* like, and prescriptive, meaning they focus on what people of a certain group *should be* like. For example, in the sex or gender context, we often encounter the expectation that women will be caring, will be nurturing. Violations of that stereotype are problematic. We also encounter expectations that black people will be violent, that they will be criminals, that they will be aggressive. And violations of those “negative” stereotypes can also cause problems just as violations of “positive” stereotypes can cause problems.

Prejudice, unlike stereotyping, is not concrete—at least, not in the sense that it can be attached to a particular idea, a particular belief or characteristic. It's more of a generalized, affective emotional response to a group. It refers to the reaction that we have, maybe the unease that we have, when we encounter people from groups that we're not familiar with. So when you're walking down the street, and you encounter somebody from a group that you're not familiar with, maybe you think, "Eh, I'm not quite comfortable with this person." You can't necessarily articulate why that's the case, and what the term “prejudice” is capturing is your inarticulate emotional response. Which is why I say that prejudice is diffuse. It's usually positive or negative in a general sense, but it's not something you can express in the form of, "I feel this way because I believe this about this person." It's just comfort or discomfort, positivity or negativity, toward social groups.

Discrimination refers to our propensity to treat some groups differently than we treat others. In a way, it's actually the behavioral response to—or behavioral manifestation of—prejudice or bias. But discrimination can also exist in the absence of prejudice or bias; all you need is a system that's set up to treat certain types of people differently than other types of people.

Discrimination, like prejudice, can be diffuse, can be vague. We obviously accept that treating certain people, or certain types of people, differently than others is ok in some circumstances. How do we decide when it's ok and when it's not? What crosses the line? The point for the moment isn't to answer that question, but just to note that discrimination has fuzzy boundaries, and any kind of differential treatment, any kind of behavioral response to a certain type of person, could in some sense be considered discrimination.

So having sketched those three concepts—stereotyping, prejudice, and discrimination—I want to start talking about implicit bias, which we hear about a lot in the post-Ferguson context. In particular, we hear a lot about how the answer to everything that's going on in the world of policing at the moment is giving police officers implicit bias training. I'm going to give an overview of implicit bias, including how it fits into our lives in general and how it's relevant to police training.

Implicit bias emerges from the idea that in terms of the way that we process information, we have two systems at work in our minds. One system operates relatively automatically, and that's the system where implicit bias originates. We also have a system that operates relatively deliberately, and that's the one where explicit bias originates. Obviously, these systems are working all the time, and the picture of how they interact is complex. But the bottom line is that when you have some kind of dilemma or question, maybe just a response to a social scenario—it could be any number of things—you have to process information in order to decide how you're going to behave. And if there is no reason for you to pay special attention to specific things involving the dilemma or social scenario, then you're going to engage in what's called automatic processing.

For example, when I go out onto Michigan Avenue and I see the rush of people coming at me, there is no reason for me to invest time or resources or to be motivated to process each of those individuals *as individuals*. Instead, I'll process them in a more generic way, like: "That's a black person." "That's a white person." "That's a woman." "That's a man." And what will automatically be activated in the course of that processing are ideas or associations that go along with those group memberships. What I won't be doing is finding out anything particular about any individual, or thinking about one person as somehow different from any other person who fits that social category.

Thus, when it comes to implicit bias, what we're suggesting is that it's a type of bias that runs through our automatic mental-processing system. It's a type of bias that does not require that we be *motivated* to think or behave in a certain way. We can experience and be affected by implicit bias without being aware that it's happening. And it's not something that happens because we want it to happen, in most cases. It's just *there*, and it's operating under the surface all the time.

In contrast to that, the deliberate system is where we find our kind of spoken—or at least articulable—beliefs and attitudes about various ideas. But that system requires motivation and cognitive resources and time and deliberation. When we look at how often in our day-to-day lives we're engaged in that conscious, motivated, deliberate processing, it's not very much, considering the amount of social stimuli we face in the world. Most of the things we respond to, even on the way to and from work, we don't have to think about consciously. We just do it. The deliberate system doesn't actually account for a lot of our minute-by-minute information-processing. It wouldn't be practical. We have to process so much information that the only way we can do it is if we mostly use our automatic cognition.

So the distinction between implicit and explicit bias comes from this dual model of cognition. And one of the things that makes implicit bias such a challenge is that it's happening beneath the surface, without our conscious recognition, and it requires effort to even be aware of it—which means it requires us to devote extra resources to monitoring our own information-processing. That doesn't usually come naturally to us, and it means we can't study or respond to implicit bias in the way we study and respond to explicit bias.

Again, explicit attitudes are the ones that we have heard of for a long time and that we feel like we have a much better grasp on. They're the attitudes we consciously, deliberately endorse. People will often—possibly for reasons of social expediency—choose not report them, but they're *able* to report them. If I ask you about your racial attitudes about a particular group, you can report what those attitudes are on a survey, for example, or on some kind of psychological measurement. If I'm doing it in that way, I'm capturing your explicit attitudes about that group. They're going to be your deliberate attitudes, things that you have, in theory, consciously decided: "I think this way about men." "I think this way about women." "I think this way about black people." "I think this way about white people." Whereas implicit attitudes are subtle and non-conscious. Usually, people can't self-report implicit attitudes. As a result, you have to develop some other kind of tool to measure implicit attitudes.

The other thing that's important about implicit attitudes is that they can be influenced by a number of different variables. They're certainly linked to our ideas that come out of our explicit attitudes, so they're linked to those consciously endorsed beliefs, but they may not always line up perfectly with those consciously endorsed beliefs. For example, I may have very strong, explicit, egalitarian attitudes or values, but I may also, when given an implicit measure or task, register some kind of bias towards certain social groups. That is not uncommon. And certainly, that's the context in which we're working, a world in which there is a lot of explicit endorsement of egalitarian values, of the value of diversity, of all these things that would suggest that we're in a post-racial society. But work on implicit bias suggests that we're not, largely because of that divergence that we often see between the explicit level and the implicit level.

Implicit attitudes, rather than being measured by self-report—by you filling out a survey or filling out a questionnaire that asks about your attitudes—are measured using other means, the most common of which is the Implicit-Association Test, or IAT. The idea is that in order to capture relatively automatic, spontaneous, less controllable attitudes, we need to use a different type of measure. We can't just ask people what they think.

I'm going to have you in just a moment go through an exercise that demonstrates what this looks like. But first, I want to point out that explicit and implicit attitudes don't just diverge: they also converge. Sometimes, they line up with one another. If you ask people about, for example, their views on various presidential candidates, if you ask them what their explicit attitudes are by giving them a questionnaire, they're going to tell you that. If you give them an implicit measure on the same topics, it will probably line up. If you ask them what they feel about a number of attitude-objects, for example, preferences for food or soda brands or something like that, you tend to get convergence between the explicit and the implicit.

Where you get a lot of divergences between the two is on controversial or sensitive topics, like views on sex or gender, or views on race. This divergence matters because explicit and implicit attitudes predict different things. So if we're looking at the relationship between an attitude and a behavior, an explicit attitude is going to predict more deliberate, more consciously decided behaviors. If I'm trying to behave in a way that is not discriminatory, my explicit attitudes are going to do a better job of predicting that. For example, if I'm trying to review a set of job applications and not take into account the race that the names on those applications suggest, my explicit attitudes are going to do a pretty good job of predicting the likelihood that I will not engage in discrimination in that setting.

Implicit attitudes, on the other hand, predict less controllable or more ambiguous or subtle responses, so they predict, for example, non-verbal interactions or non-verbal communication. How likely am I to sit next to a black person on the bus? How close am I willing to be to people from different groups that I'm not familiar with? Those are things that we don't necessarily consciously decide. They're kind of related to this idea that we feel these general levels of discomfort or comfort in relation to certain people. And that's more closely tied to implicit attitudes than to explicit attitudes.

Going back to the resume example, if I have no conscious goal not to be discriminatory, if I'm reviewing two resumes or applications with names that suggest one white applicant and one black applicant and I don't have a conscious desire to not discriminate, that's when implicit bias can play a role. I may be drawn to the white candidate over the black candidate because they just seem like they're a better fit. I can't articulate exactly why that's the case, but I'm drawn to that person in some way.

It's not so much that explicit and implicit attitudes always produce different responses, but that the way that you get to the responses may differ depending on whether you're taking the explicit route or the implicit route. Consequently, the way that you address problems has to be different, depending on whether you're going after the explicit or the implicit.

And I'll say that in the context of bias in policing, we're certainly not in a world where the only type of bias we're concerned about or should be concerned about is implicit; there are certainly lots of explicit biases at play alongside implicit biases. You hear a lot of discussion about implicit bias in policing, and implicit bias training as being the solution to the problems that we see, but we have just as many problems in society and policing generally that have to do with explicit biases. In fact, post-Obama, rather than seeing a post-racial world, we see an increase in expressions of explicit bias. We see an increase in the kind of blatant, overt discriminatory or prejudicial behaviors that we had hoped were long gone, or at least going away, in our society.

So it's not that implicit bias is the only thing we have to be concerned with, it's more that even measuring it and bringing it into the open poses challenges that we're working on surmounting. As I mentioned, the main measurement of implicit bias is the Implicit Association Test, the IAT. This is an evaluative test that's meant to measure the general positivity or negativity that you feel toward certain social groups or attitude objects. It can measure any implicit bias or association, not just those associated with race. People often think of it as being closely connected to measuring racial bias, but you can use the IAT to measure people's preferences for Coke versus Pepsi, for certain presidential candidates; certainly, other social categories beyond race.

One thing to know is that it's not considered, at least by researchers, to be individually diagnostic, which means it's not a test or measurement that you can administer and then label somebody as being racist or not racist, biased or not biased. The information that it gives us is useful in the aggregate. It's information that enables us to see, generally, that people move in one direction or another on these results. It can tell us about how people tend to think, not necessarily how any one individual thinks. There's some dispute about this with regard to the difference between what researchers get out of IAT results and what laypeople get out of IAT results. In practice, most people want to be able to have an individually diagnostic tool, and that can lead to a discrepancy in how laypeople and researchers view the IAT. This is compounded or amplified by the fact that the test is

administered on an individual basis. So it's easy to see ways in which we'd assume it was individually diagnostic. But researchers caution against using it for those reasons.

I wanted to bring up those points because what I'm going to do now is actually give you an opportunity to go through a modified version of an IAT. And I assume that at least some people in the room have already done one of these, so bear with me. But for those of you who may not be familiar, this will give you a general sense of how implicit tasks tend to operate. After this, I'll talk a little bit about other variations of the task that apply specifically to situations that are invoked by the topic of today's symposium: implicit bias in policing.

So you're going to complete a task in which you will sort words and pictures into categories as quickly as possible. All you have to decide is whether a particular word is positive or negative, or whether a face that you see is a black face or a white face.

Here are some examples of the words that you'll see. Positive words include things like joy, love, peace, et cetera. Negative words include things like agony, terrible, horrible. So it will be clear—it *should* be clear—which are the good words and which are the bad words. You'll also see faces of black people or white people. Again, it should be clear whether the face is a black face or a white face.

Now I need you to take things out of your laps. And I need you to put one hand on your left thigh, and one hand on your right thigh.

What you're going to see in a moment is two labels at the top of the screen. And those labels are going to indicate whether you slap your left thigh or your right thigh. Remember, all you're doing is categorizing. All you're doing is deciding, Is this a black face or a white face? Is this a good word or a bad word?

You want to go as fast as possible. You want to be as accurate as possible. But if you make a mistake, it's okay. Don't worry about it. If you do realize that you've made a mistake, simply just correct to the other side. So if you realize you accidentally slapped the wrong thigh, slap the other thigh before moving on.

Does everyone feel ready at this point? So everyone is nervous. It's okay. Remember, this is not individually diagnostic.

Here is our first set of categories. You have white American on your left side. You have black American on your right side. So one hand on each knee. Faces are going to appear. Simply slap whichever thigh corresponds to the face that you see. Be as accurate as possible. Do it as quickly as possible, okay?

Ready? Okay. So we're going to change the categories. So now, you have good on one side, bad on one side. And this time, you're going to see those words that we were talking about. So you see the word, and you just have to say is this a good word or a bad word. If it's a good word, you slap your left thigh. If it's a bad word, you slap your right thigh. Okay?

You got faster that time. I heard that.

So now, we're going to pair these two things together. So it's the same side as when you did them separately, but you're going to see a mixture of words and faces. If you see a white face or a good word, you're going to slap your left thigh. If you see a black face or a bad word, you're going to slap your right thigh. So the same words, same faces that you were seeing before, same side that you were slapping on before. It's just that we've now paired the words and the faces on each side, okay?

Now, we're going to flip the sides. So now, you're going to see the same faces again, black and white faces. But now, you slap your left side if you see a black face. If you see a white face, you slap the right side. So again, the same faces that you were seeing before. We're simply switching sides.

Now we're going to add the good and the bad back in. So slap your left thigh if you see a black face or a good word. Slap your right thigh if you see a white face or a bad word. I know this is hard. Bear with me.

So that's the Implicit-Association Test (IAT). And what that test is measuring is how strongly you associate blackness with goodness or badness, or whiteness with goodness or badness. You can hear the pauses that people have when you pair the counter-stereotypical responses. That's because we have strong associations, not just on the basis of social targets, but on the basis of general ideas of blackness and badness, whiteness and goodness, which means that when you flip those pairings, it's harder to make those responses.

What that test illustrates is how much harder it is to make counter-stereotypical pairs than it is to make stereotypical pairs. And what you find in these measurements is that people generally have trouble with the counter-stereotypical pairs: matching white with bad and matching black with good. The idea is that the strength of the association predicts whether you're likely to engage in behaviors that are informed by implicit bias, and that produce disparities on the basis of implicit bias.

There is a lot of debate about whether that test, or any tests of this variety, reflect *personal beliefs* or *cultural influences*, like stereotypes. That debate is ongoing. It's been ongoing for basically as long as the IAT and implicit bias has been researched. But it doesn't matter, because what we do see consistently is that the strength of those associations, the quickness with which you slap the correct knee, can predict a variety of behaviors, regardless of whether your responses come from cultural influences or from personal attitudes. You can see it in hiring/firing/promotion decisions in the workplace. You can see it in comfort, closeness with others, allocations of various resources and opportunities, and certainly in reactions that police have to, say, white suspects versus black suspects.

Now I want to mention some additional tasks that aren't exactly like the IAT but are meant to capture the sort of automatic, fast, spontaneous responses that are particularly relevant to policing.

First is the weapons identification task. In the weapons identification task, what you have is a presentation of faces. Think of the faces you were just looking at in the Implicit-Association Test. You have a presentation of a white face or a black face. And in the weapons identification task, after you see that white face or black face, you see either a weapon or a tool. What the researchers look at is whether the color of the face affects identification of the next image as a tool or as a weapon. And what they find is that when a black face precedes a weapon, you get faster identification of that weapon. When a black face precedes a tool, you get more errors in which that tool is identified as a weapon.

Second is the sheer bias task, which many of you may have heard of. In the sheer bias task, police officers, or study participants who include police officers, are presented with scenes containing a potential suspect. Sometimes the suspect is white. Sometimes the suspect is black. Sometimes they're standing, sometimes they're crouching. The scenes include various settings out in the world, some urban, some not. And the person—the potential suspect—is holding in one of their hands either an innocuous item, like a cell

phone or a soda can, or a weapon. Study participants—again, including police officers—are asked to decide whether they should shoot this person. The instructions are to shoot people who have weapons. And what you find is that, for one thing, police officers do significantly better than laypeople at this, because they've received some relevant training and it actually makes a difference. But you also find that in general, people are faster to shoot black men with weapons than they are to shoot white men. If anything, there's an inhibition effect surrounding white suspects with weapons: people are particularly slow to shoot them. There are also more errors, in the sense that there are more shootings of black men who don't have weapons than white men who don't have weapons.

These tasks are operating on the same basic principles as the Implicit Association Test, but they're doing it in a way that is much more aligned with policing.

Additionally, there are other processes at play in the policing context that aren't exactly implicit bias, but are certainly relevant to our discussion. One is dehumanization. Dehumanization refers to associating social groups with non-humans, like animals. There is research that shows that increased dehumanization predicts greater use of force by police. For example, the Goff study I refer to here looks at the use of force against black children. Dehumanization can attach to any number of social groups in the policing context—particularly men of color, but certainly all people of color, people who are convicted of crimes more generally.

Another important threat or bias that might be at play in the policing context is stereotype threat. Stereotype threat refers to the idea that fear of reinforcing a negative stereotype about your own social group can lead you to underperform in the domain in which you're working. Thus, in the policing context, officers' fear of being perceived as racist can put them under threat, which interferes with their ability to actually behave in a non-biased way.

Similarly, masculinity threat is important in the context of policing. Within police culture, a perceived threat to one's masculinity can lead to overcompensation in the form of force and violence. This type of overcompensation can even help explain the use of force and violence by women and LGBT police officers, because the highly masculine culture of policing can make people susceptible to masculinity threat even when they don't necessarily think of masculinity as a big part of their personal identity.

We're focusing on policing today, but I want to mention that there are other areas of the criminal justice system in which we see the same or similar biases having an effect. One example is sentencing. People who look "more black" tend to be sentenced more harshly. A more stereotypically black-looking man, like the one on the right, is more likely to receive the death penalty than the one on the left. And that's even when you control for basically every other important factor that should affect sentencing.

The steps we're taking to mitigate and eventually solve these problems include increasing awareness (which is why we're all here today), increasing accountability, increasing transparency, and incorporating diverse perspectives in our institutions. And I'll say that when it comes to diversity, it's not as simple as just hiring diverse people—you need to *do* something with the diversity. Which is true, in some ways, of all of this.

I'll take questions now.

VOICE: Hi. I really appreciated your talk. I had a quick question about the training that police receive: is it transparent? Can laypeople or citizens see the training that police get and contribute to that training?

PROFESSOR PEERY: We don't have a lot of specific information about police training. And that's in part because training differs by police department. There is no universal training for police. One of the big pieces that I know the most about, and which has come into use more and more, is simulation training, where you engage with suspects in a virtual reality setting. And you have to make those shoot or not-shoot decisions.

There is some research by Phil Goff into those simulation trainings, and that's part of where he came up with the idea of masculinity threat and stereotype threat being relevant factors alongside to implicit bias because. For example, he found that in those simulation trainings, it wasn't implicit bias that predicted officers' behavior, it was whether or not they were a little worried that someone would perceive them as not being masculine enough, or worried that someone would perceive them as being racist.

That's simulation training. Police go through lots of other forms of training, and we don't know a whole lot of specifics when it comes to those other forms.

VOICE: Hi. Thank you for your presentation. My colleague and I are federal defenders, and so we're wondering: what are some of the ways in which you've seen these kinds of concepts implemented at the podium, for example, in a criminal case? How can we talk about implicit bias and stereotype threat with judges?

PROFESSOR PEERY: It's certainly hard see how you get this information into the courtroom, and it might be more promising to think about keeping individual judges informed of these ideas in general. Judges go through a lot of training, and I think the response thus far has been to try to increase, over time, the extent to which we train judges and attorneys about implicit bias, about stereotype threat, about how these things can affect their decision making. There are some judges who have decided on their own accord to, for example, give anti-bias jury instructions, or who have made commitments to reminding themselves, or have taken pledges to not be biased. But again, that's a matter of individual judges deciding how to run their courtrooms.

It is much trickier to think about how to introduce these ideas as evidence in a case. Certainly, in the Trayvon Martin case, you have the attorneys wanting to introduce evidence of race and the role that that played, and the judge rejected that. So I think you have to try, but I'm not sure at this point what would be truly convincing for judges. And in part, that's because the implicit bias evidence or the implicit bias information doesn't quite match up with our ideas of evidence and what constitutes evidence or what is deemed relevant evidence in a court case versus how we see it as being relevant maybe outside of that sphere. So I'm not sure that I have a good answer for you at this point.

VOICE: Hi. I'm wondering if statistically, you've been able to take a look at regionality with the sentences. You mentioned the sentencing of black individuals versus white individuals, and I assume for the same crimes. But is there any difference between circuit courts and federal courts in the specific area and regional differences?

PROFESSOR PEERY: I'm not sure there has been research at the kind of level of the court. Even sentencing information of that variety can be hard to get. The studies that have been able to do that have gone through extensive work to pull a lot of records where you can actually get images of the individuals. After the fact, they're doing basically an archival analysis, paired with the kind of rating of these people's appearances. And that can be very hard to compile. Just like a lot of other data in the criminal justice system, it's not always very well kept. And it's also not something that people can always get access to.

There has been quite a bit of research on the death penalty in federal sentencing and the effect of federal sentencing guidelines in particular. Which might suggest that that data is easier to get than the kind of lower state court or district court level—more at the appellate level or in capital cases, where there tends to be more of a record than exists in your run-of-the mill criminal case.

VOICE: Do you think the weapon identification test and the shooter bias test would be useful in deciding whether to hire people as police officers?

PROFESSOR PEERY: That's been proposed. The pushback actually comes from researchers who have developed the tasks, who maintain that these are not individually diagnostic tools. The research, even after multiple decades, has not advanced to the point where we feel comfortable saying that any of these tools should be used to make some kind of individually diagnostic determination.

Now, that's not to say it's not going to happen. I think the concern would be that it might actually disqualify too many people. So using that as a criterion for qualifying people to be police officers might be problematic. One of the reasons is if you're looking at it through the lens of, say, anti-black bias or pro-white bias, the default is for us as a society to have the pro-white/anti-black bias. There's variation among different groups, but largely, it's skewed toward that, which means a lot of people will fail on that dimension. And I think that would be the concern about using the test in that way.

VOICE: Hey. I am with First Defense Legal Aid. If you are a lawyer, you can actually take a six-hour shift with us. And what we do is we go directly to the police station when someone is arrested or detained. While we're trying to work on changing judges' minds or police officers' minds, we can also be working on upholding civil rights and making sure those biases don't even make it to the courtroom by having a lawyer at the station. You can come and leave your name. And like I said—I will tell you more about it—you could get a six-hour shift once a month and go directly to the station. If you are just an organizer or a nonlegal worker, we also do Know Your Rights workshops. And I have cards that I also can hand out. So you can see me after the break.

PROFESSOR PEERY: That highlights an important point: in terms of thinking about how you increase accountability and transparency, it's not just about the individual, it's about the system. Going after individual bias is really hard, particularly when you're talking about implicit bias. What's more likely to be influential is systemic change: changing processes, changing policies and procedures, to inoculate against bias. It's not a problem to try to go after bias directly. But a lot of times, that's not very effective. A lot of diversity training, et cetera, is not very effective. There are instances where it's more harmful than it is helpful.

Whereas changing policies and procedures to create systemic effects can undo some of individual bias.

VOICE: I do juvenile records expungement. I often see cases where children are charged as adults. And it seems to me there was a study about the perception that African-American children were older or more responsible than they actually were. And I'm just wondering if we can incorporate that. I don't remember the study.

PROFESSOR PEERY: Yes. So there have been a couple of studies at this point showing that the implicit bias people have regarding black people attaches even to children. Even though we might think it applies predominantly to men or adults, we also see evidence that it applies to children. So the stereotype or the bias that black people are violent or criminal or aggressive attaches to children.

And when it does, it means that we perceive those children as being older than they are. On average, black children are perceived to be four years older than their actual age. And with that comes greater acceptance that they are morally responsible, that they have more individual liability and responsibility for their actions. They're not granted the same things we grant other youth: impulsiveness or the stupidity of being young. Going out and doing dumb things because you feel peer pressure and whatnot. We skip over that and assume that they're more adult than actually they are, that they're more developed cognitively, morally, emotionally, et cetera. Which makes it okay in our minds to punish them more harshly—or at least in the system's mind. Because, you know, the criminal legal system is built on the idea that we want to punish people who are more morally responsible or more culpable in some way. So that is a finding that is out there.

I will be around, so if you have other questions, certainly let me know.