



Opinion Science Podcast

Hosted by Andy Luttrell

Episode #16: Implicit Bias with Mahzarin Banaji

July 27th, 2020

Web: <http://opinionsciencepodcast.com/>

Twitter: [@OpinionSciPod](https://twitter.com/OpinionSciPod)

Facebook: [OpinionSciPod](https://www.facebook.com/OpinionSciPod)

Andy Luttrell:

Look, I don't know who you are, but I feel like there's a good chance you've heard of something called implicit bias. Usually we talk about this as a kind of prejudice: racism, sexism, ageism, that lives beneath the surface of our conscious minds. The idea is that even when people openly embrace equality and vehemently deny that they act with bias, there can still be prejudices floating around in their heads. Or as the breakout Broadway hit, Avenue Q, puts it, "Everybody's a little racist." A couple years ago, Starbucks got into some hot water, and I swear, I didn't realize that was a pun when I wrote it, but I decided to keep it in, didn't I? Anyhow, in 2018, two Black men were waiting for a friend to show up at a Philadelphia Starbucks, but because they hadn't ordered anything and were just sitting there... I don't know, menacingly? Starbucks employees called 911.

This sparked lots of conversation about how people can make assumptions of others based on race, even if they don't realize they're doing it. CNN put it pretty simply with the headline, "What the Starbucks Incident Tells Us About Implicit Bias." And how to combat this sneaky, hidden type of bias? Well, if you're Starbucks, you shut down all of your stores for a day and administer implicit bias training to your employees. And it's not just Starbucks. Implicit bias training has become a big industry as we seek new ways to cast out these unwanted assumptions of other people. But where did this idea of implicit bias come from? How can we measure biases that people don't know they have, or at least are unwilling to endorse openly? And is implicit bias training really enough to address the inequities in our society?

You're listening to Opinion Science, the show about our opinions, where they come from, and how they change. I'm Andy Luttrell, and this week I had the distinct pleasure of talking with Dr. Mahzarin Banaji. Along with her colleague, Tony Greenwald, and their long list of collaborators, Mahzarin pioneered the research in social psychology on implicit bias, a body of work going back to the mid-1990s. And a few years ago, she co-wrote the book *Blindspot: Hidden Biases of Good People*. I talked to Mahzarin about her earliest days as a social psychologist and how she got wrapped up in studying implicit bias in the first place, and she also shared new research with her student, Tessa Charlesworth, about how implicit biases can change and what that means for our ability to make real social progress. This one's a little longer than other episodes, but I hope you enjoy our conversation as much as I did.

Mahzarin Banaji:

Okay, so let's just start with some definitional things to get them out of the way. So, I would say that implicit bias is actually a term that was coined by scientists to refer to behaviors that tilt away from equality. Something that is not neutral, even when neutrality is what you desire and neutrality is what you believe you've actually demonstrated, right?

So, the bias part comes in when we speak about a behavior that tilts away from equality, and the implicit parts is embedded in the fact that the person who is doing the behaving, and maybe even the person who is receiving the behaving believe that that behavior was enacted in an equal way. So, implicit bias then is most clearly seen when you have let's say two people before you, both equally qualified, both equally trustworthy and safe, and your preference systematically leads you to select one over the other. A white applicant over a Black applicant, or a Black applicant now is equal to a white applicant when the white applicant's resume is identical in every way, except has a comment on it that says that that white applicant is also a felon, and now when they get treated equally, we can say it's a deviation from neutrality. When doctors prescribe pain medication differently to people from different racial groups, and the pain data are the most clear because it's in every disease, in every geographic region of the country that Black Americans get prescribed lower amounts of pain killers for reporting the same level of pain on a scale.

And when this happens with no awareness on the part of the doctor, and I have good reason to believe that these doctors are like me, that they have not a shred of bias. If anything, I think they're thinking about it and wondering if they're overcorrecting too much. That's their usual worry and the data don't show any overcorrection, they just show the same old thing. So, if you're not aware that is a person's skin tone, or their facial features, or their dress, or their accent, or their religion, or nationality, or age, or gender that's causing the bias, and yet when that feature does intrude into decision making, we say that that bias is implicit.

Andy Luttrell:

It sounds... Is implicit bias a special case of an implicit attitude? So, your definition is very clear on the bias side, but my impression is that some of this arose out of just an interest in understanding just preferences or attitudes in general that might exist at this same kind of level.

Mahzarin Banaji:

So, people have come to this topic in a few different ways, and I can tell you what my path was, and that may clarify things. When I was at Ohio State, I was, even at that time, mostly interested in human memory. Almost entirely in human memory. Master's thesis on how the self ends up being a memory system to which information becomes attached, that information achieves a different status in memory. It's better remembered, it's less forgotten, it has more coherence and all of that. My dissertation was just an extension of that, but into the question of affect. To what extent is information that is affectively positive or negative remembered better compared to neutral, or is it? So, it was really simple, two hypotheses. In those days, we would say, "Is there an intensity effect in memory such that anything that's equally deviant from zero in the positive or negative direction will be well remembered compared to what's neutral? Or is there still this Freudian notion of repression, where negative stuff gets repressed and forgotten?"

So, again, so never collected any data ever with any dependent variable measure other than memory. All kinds of memory measures. And then when I got to Yale as an assistant professor is when the implicit memory revolution was happening with cognitive psychologists who were studying memory discovering that something interesting is going on, that the kind of memory we'd been studying as psychologists for over 100 years was what we would today call explicit memory. You tell somebody, "We gave you a list of words earlier. Can you remember all those words?" We can say to them, "What did you eat for breakfast? What happened to you when you were five years old and you went to a fair and got lost?" We can ask people to do that and each of those tests of memory is the same in the sense that you're requiring the subject to go into their minds to some past moment in their history and to pull information out.

And today, we would say that's one kind of memory, but for the first 100 years, psychology assumed it was the only kind of memory. And then as we saw, first from data from patients with amnesia, that it appeared that even though they had nothing like the kind of memory you and I have, where a few minutes after an incident, they would not remember the incident ever occurring, that even though they couldn't remember in the classic sense, do you remember the joke I told you? "No, did you tell me a joke? I don't even know who you are," they would say. But on each subsequent telling of the joke, the joke would seem less and less funny to them. Even though there was no conscious recollection of the joke.

Anecdotes like that, and then actual tests that showed that amnesic patients had a form of memory that we had missed, that we might call implicit memory. And how did it work? You could give them a list of words at time one. At time two, you would ask them to remember those words, and as I said, they would not only not remember any of the words, they wouldn't even remember that you had them do such a learning task. However, if you gave them a letter and said, "Fill out the rest of this word. It's a word that starts with an M and it's five letters long," they would be much more likely to complete it with the word motel if motel appeared on the list before. So, these kinds of data were just stunning, and it led to I would say a revolution in our understanding of human memory, that memory comes in many different forms, that there are certain kinds of memory that people... And Larry Jacoby was a very major influence on me, and he had done some beautiful experiments.

And all I did is I thought I would just replicate Larry Jacoby's experiment as an assistant professor just thrashing around and not knowing what I wanted to study, except that I knew I was interested in memory and I would go to a journal club that met every Friday on memory, and I presented Larry's paper on the false fame bias, in which he showed that if you'd seen a name before, a name like Simon Weisdorf, at a later time you'll be more likely to mistakenly identify that name as a famous name.

So, I, in those days, this is about still in the late 90s, or sorry, in the late 80s, and I write to him and I say... You know, we used to send postcards in those days, so I'd send a postcard and the postcard would say, "Dear Dr. Jacoby, I loved your paper. Would you please be able to send me any materials that I might use to replicate this?" And then you would wait, and wait, and wait, until you saw a fat manilla folder in your mailbox, and you would pounce on it, because that would be the stimulus materials that would have arrived. And as I looked at Larry's stimulus materials, I

noticed that all the names he had used, most of the names, like 90-something percent were male names.

So, as I say, you write another postcard. Dear Dr. Jacoby, you only sent me half your stimulus, could you send me the remainder of it? And if you're lucky, you'd get a call. A gruff voice at the other end saying, "Well, you know, women aren't famous, so we didn't use very many women's names. Bye-bye." And I remember that when he said that to me, I remember thinking, "Yeah, he's right. There are not many famous women, and so it would be best to just leave them to be a few." And this is where I would say to a young person entering the field who thinks that science is pristine and not affected by anything that is personal, and I would say to them to pause, because I didn't consciously at all think that what Larry had done was wrong. I felt like, "Yep, I can easily just replicate the study with his stimuli."

But I didn't. I actually spent an entire summer generating names of famous women, and I just did the study again with both kinds of names in equal numbers, with the expectation that I would just replicate the Jacoby result with just a broader set of stimuli. And the result replicated, but for male names. After the first study showed that result, I assumed that all the subjects were actually doing exactly what Larry said. They know there are few women who are famous, and so when they see a name and it's the name of a woman, they say, "What's the likelihood that she'd be famous? I'm not gonna circle it." And that's what I thought was going on.

So, starting with experiment two, I began to in the debriefing ask the simple question, "What did you use in your decision?" And people would come up with lots of crazy ideas, but they actually nailed the Jacoby hypothesis. They said, "I can imagine that if you had shown me a name the day before, that I might be confused today about why is it feeling..." What Jacoby called this was perceptual fluency, when you see a name like Sebastian Weisdorf again, it has a degree of perceptual fluency. And some of them would tell me that. They would actually generate the Jacoby hypothesis. But not one said it could be the gender of the name, so then at the end of their telling me their hypotheses, I would say, "Do you think the gender of the name affected your decision?" I started to code this.

400 subjects later, not one agreed that the gender of the name had mattered. And I knew... I didn't know what it was gonna become, but I knew in that moment that this result was more important than anything on memory I could study. And I think that was the moment when I said, "It's time to become a real social psychologist."

Andy Luttrell:

Yeah, so from there, what does the road look like to get to... to circle back to what implicit bias is and how it's different from explicit bias, these sort of rumblings of memory processes, where have they taken us?

Mahzarin Banaji:

Yeah. It's a perfect question, because remember I said to you that you come to this work along a few different paths. I could tell the story the way I just did, but alongside that story is something else. And really, issues of discrimination and equality were not really being talked about. Remember that in the 1970s, social psychology went through something called the crisis, and the

crisis was that in the post-war generation and in the people trained by them, they had been interested in things like authoritarian people, and personalities. Adorno and all of those guys who had come from Germany were very struck by what had happened in Germany, and it bled into their work. But in the 1970s, there was what was called the crisis, and the crisis was must we be studying these questions of relevance, or can we be scientists just studying something for the hell of it?

And the field kind of cracked open along those lines, and when I came to Ohio State, clearly the side that said, “No, you don’t have to study anything of relevance in the real world. Scientists study what are important problems.” And so, by the time I came, no mention of anything like discrimination, inequality, race, gender, none of those were being talked about. However, we could see something. When we took a course, you could see that the Katz and Braly data, collected in 1933, showing that people very strongly believe that Italians were creative, musical, lazy, and dirty, that that kind of responding was simply not showing up 20 years later, and then another 20 years later, there is something called the Princeton Trilogy, the set of sort of three papers, each appearing 20 years apart, showing that the stereotypes of 20 years ago were simply not appearing.

So, if you look at the data, and in our book, *Blindspot*, we have an appendix that is devoted entirely to this issue of showing what has happened to race attitudes over time, what you saw is a significant drop-off in anti-Black bias. So, that was happening, but if you had even two neurons in your brain, you could see that in the real world, there was something happening that did not seem to match with this fall off in reporting bias. In the real world, there were differences in who could get a bank loan. In the real world, people’s health outcomes were different. In the real world, there were mortality differences. In the real world, you could see that there were differences in who worked in what job.

So, how to line this up and resolve this paradox that on the one hand, people are just explicitly telling you that they’re not biased, and you have no reason really... We’re not silly people who think people are just lying to us. We could tell that they were expressing the truth of how they felt. And yet that’s what the world looked like, so why wasn’t there a parallel shift, so that what we were measuring in the world was looking like it was following the shift in attitudes and stereotypes falling off from negative to neutral?

And that itself was also interesting, and that became a methodological interest. And for me, it has always been more interesting to think of implicit measures not at all as a way of getting at the truth, but as a way of getting to an aspect of our minds that we are not in contact with consciously, and it’s the same as with memory. You don’t go in and say, “What did you eat for breakfast?” You just give somebody the word O and they’ll fill it in with omelet if that’s what they saw on the menu in the morning. That was the parallel. So, there was a very clear parallel between, “Okay, one of our big concepts, memory, is cracking. It’s showing that there are at least two kinds and we are gonna do the same for attitude.”

And I always said that, that there’s nothing original in our work. That all we did is we took what was going on with memory and we simply asked the question, “Is it also true of our attitudes? Is it possible that they come in a couple of different forms? The conscious kind and then the less conscious kind?”

Andy Luttrell:

Yeah, so to take that, what I was thinking when you were talking about the memory example and amnesics in particular, like we have very clear statements about awareness and consciousness, and certainly over the years, implicit bias has been characterized as unconscious biases. These attitudes, preferences, biases that people just do not at all realize they have. But as you're well aware, that claim has been debated over the time that this research is being done. So, I'm curious at this point, based on the evidence that we do have, how strong a case do you think that there is that these are attitudes that exist outside people's awareness?

Mahzarin Banaji:

Well, there is some truth to those 400 people, who all told me that they were not using gender. Not one. So, I would say something's going on. If it were the case that these were all within conscious awareness, we wouldn't get the hate mail we do, okay? So, all I can tell you is that that is the signature result, that responses on two different measures don't line up. That's the signature result, so we have to agree with that. However, I think your question is a very good one, because in the old days, I think we were mistaken. I think we thought these were two separate systems with not much in common, that these were really, just like with memory, they were separate memory systems. That this is separate systems in the brain, that one responded to the conscious probe, the other to the less conscious one, and they didn't speak to each other. And that's largely been shown to be false.

What do I mean by that? Very early on, this was research I did with Wil Cunningham at Yale University in which we did, I think, sort of the first... In those days, people used Ns of 30, and we went big, and we said, "We're gonna pick five explicit attitude groups. For those same, we'll have five different relevant implicit measures. We'll put in every explicit measure for each one of those that we can find that people have used, and we're gonna get people to come in and take them multiple times." All of that. So, we do all of that, and the result of that paper that somehow doesn't get reported and never even I think led people like myself to say, "Whoa! What are we talking about? Look at the data."

And what we showed is that if you look at the latent structure... Let's say these are age attitudes, foreigner attitude, race attitudes, whatever it was. I forget the five, but whatever they were, we showed that the cluster of implicit attitudes hung together. The explicit hung together even more beautifully. And the correlation between the two was .5. But it was always a two factor solution that fit best, never a one factor solution. And I think that that, for me, still encapsulates what I think is the truth of their relationship. It's the same universe in which we get our experiences, they get filtered through different senses. The stuff that we learn, parents will say to us, "You know, we never teach our child stereotypes and we don't know where they're getting them from." I say, "Did you ever say to your child, as a father, did you ever say go get mommy from the kitchen? As a mommy, did you ever say to your child, daddy's gonna be late at the office today?" That's teaching them the stereotype. Kitchen and office with male and female.

And those data are getting collected somewhere in the same way as a rat would learn that when it makes a left turn it gets a pellet, but not if it makes a right turn. And one of the arguments that people made very early on, I would say largely the psychologists who studied what was called in

those days prejudice, they were the most outraged at us, right? Because we're doing two things. They're worried that we might displace them entirely, because now everybody's gonna want to do implicit measures, and what's gonna happen to their poor surveys. We had no intention of ever going in that direction. To us, the survey data are hugely important. They're telling us about an equally important part of ourselves.

And so, we never made the case that explicit anything was less important, but that's what we were perceived as saying, and sure, we were more excited about implicit measures. They were newer. You had to pay attention to them. There was something that was causing us to just spin. And then imagine you're a person, you're a white liberal social psychologist, and for decades you've been talking about those people south of the Mason-Dixon line, who are the bad people, and now you, on the IAT, are showing something very similar to those people south of the Mason-Dixon line. So, for a variety of reasons, the reaction to what we did was, "Well, this is not an attitude."

And this may be of interest to you as a young scientist. I was just so amazed by these, because remember that we were not only not the first to talk about implicit anything. We were not even the first to talk about the semantic priming type results, which John Bargh and Russ Fazio had reported on. You place a word on a screen, have it go away very quickly, another word appears on the screen, and you must make a decision on it. And lo and behold, if those two words were not semantically at all related, but affectively related, that word one and word two had a similar affective tone, mildly positive, neutral, that your response to that second word was being affected by the affect of that previous prime. And they called it an attitude, and they called it an automatic attitude, and they had gone at each other and had debates about this and so on, but we all knew this work. We taught it in our classes. Nobody said to them, "This is not an attitude."

But when the IAT came along, oh my God. The first couple of years was all about, "This is not an attitude. This is not an attitude. This is not an attitude." Well, what is it? Well, it's something in the culture. Well, yes, of course it's coming from there, but it's now in your head. But it was very hard for people to engage with that way. And by the way, I'm talking to you in a way I've never talked to any other podcaster, because this is... These are the inside stories of what went on in those days, and about that time, my friend Liz Phelps, who was a neuroscientist and in sort of that first generation of psychologists doing the neuroscience of memory, said to me, "I think we should do an imaging study using IAT outside the magnet and then collect some data inside the magnet and see if there's any kind of correlation between the two."

And I said, "I don't need a fancy machine. My test is giving me beautiful data. The statistical effects are so enormous, as somebody said, you don't need a computer to measure this. A sun dial is enough tell you that there is an effect here." So, why would I do a study with an N of 12 people in a magnet when I have beautiful data from larger numbers of people now. This was by this time, we have the web and so on.

Andy Luttrell:

And it cost quite a bit less.

Mahzarin Banaji:

Yeah. And it cost nothing. They're coming for free. So, I said to her no, and then two years later, I kind of went crawling back on hands and knees and I said, "They're not believing that it's an attitude. I think we should do that study." So, our study was very simple. Lie in the magnet. You're just doing some little tasks, but you'll see Black and you'll see Black and white faces popping up, and all we're doing is measuring amygdala activation to Black relative to white faces. And outside the magnet, we've given them a race IAT and a modern racism scale.

But here's another interesting feature about how science progresses. In those earliest days, we didn't analyze our own imaging data. Liz didn't analyze her own imaging data. She would have to have on her team somebody called an MR physicist, who would usually be in a medical school, and they would be running the magnet, and then you would pay them to do the analysis. And so, they said to us, "Well, we looked at the data. There is no difference in amygdala activation to Black or white faces." And we said, "Yeah, yeah. We know that, but we're not interested in that. We want to know if there is a correlation between IAT performance and amygdala activation." And they said, "We don't know how to do a correlation with imaging data."

So, I came back to my office and I said to Wil Cunningham, "There's a field. It's called neuroscience. Do you want to be a neuroscientist?" And I'm not sure that he responded by saying, "Yes, I'd love to." We'd been doing great behavioral work. He was a methodologist. He cared a lot about doing studies with the best designs and all of that. And I said, "Go over to the medical school and tell them not only can they do correlations, but teach them some factor analysis. That'll blow their mind." So, he goes over there and that's his introduction, and he becomes an author on this paper because as soon as they looked at the correlation, oh my God. He just called me on the phone from the medical school, "Mahzarin!" Just screaming at me. "Mahzarin! People higher in race bias on the IAT are showing greater amygdala activation to Black relative to white faces and the correlation is like .8!"

Andy Luttrell:

Can you describe what it is about the amygdala? Why was that where you were looking?

Mahzarin Banaji:

So, we picked it because of all the regions that we... Remember, our understanding of the brain was itself quite minimal when it came to which regions are involved in which psychological processes. But the one that was really beautifully mapped out because of the animal models that we had was the amygdala, right? And we knew that the amygdala is involved in something that we might call fear conditioning. So, the amygdala just actually, because of the fear part, because if we were doing gender, we would never have picked the amygdala, so race turned out to be very useful because we already knew what the amygdala should do. We knew it had been involved in fear learning. And if there is this correlation between amygdala activation and IAT, then we've got to ask ourselves, then what the hell is it?

And all I can tell you is that the IAT data were beautiful, much more robust. I think they can predict things which I don't think the imaging data could do. But all of a sudden, that question, it's not an attitude, why do you call it an attitude, just went away. Now, I don't know whether it went away, and it would have gone away anyway, or whether this paper had something to do with it, but once

we showed that that behavior on the IAT is correlated with amygdala activation, that argument just died. Yeah.

Andy Luttrell:

What's exciting about that story, too, is we're talking very early days of what became social neuroscience, right? 2001, this paper, I think?

Mahzarin Banaji:

Yeah, but we did it much earlier. We did it in the late 1990s, I think '99 or so.

Andy Luttrell:

Wow.

Mahzarin Banaji:

And there is another story that... I mean, I won't mention the journal, but a very, very prominent journal for any scientist accepted the paper and then rejected it right before publication time. And we were never told exactly what happened, but we were at the stage of signing copyright forms, and we were told that there is a panel that looks at... This was a proforma thing, but the panel of somebody who sits above the editor just looks at it to see that everything looks okay, and that group said, "This is too controversial." So, we just took the paper to the Journal of Cognitive Neuroscience and we published it there. And after that, quickly there were other people who replicated the result, and in fact, Wil went on to do a beautiful paper in which he showed that people for whom dorsal actual PFC activation was greater, now we're getting outside the realm of subcortical regions and looking at prefrontal cortex, that for those for whom what we might call the region that can exert control is more active also showed less amygdala activation.

But that was itself, I think, an interesting possibility, that for whatever reason, people who are able or have learned to exert conscious control could tamp down the default fear reaction. Yeah.

Andy Luttrell:

You mentioned that this fMRI study was very much related to correlations between IAT scores and brain activity, but we don't have to go super deep into it, but for those who may not be aware or understanding of what the IAT is, could you do sort of a brief overview of what that means?

Mahzarin Banaji:

Oh, yes. So, what is the IAT? Maybe I should just tell it to you based on how I took the test for the very first time in 1994. Tony Greenwald sent me a program and didn't say very much about what to do, but we used to exchange different protocols and try out different tasks, because we had published... In publishing the psych review paper, the implicit social cognition review paper, that paper ended by saying we're trying to make a case for this concept by relying on effects already existing in the field. This is a theoretical argument for this idea. But really, what's gonna be needed is a method.

So, Tony more than me, I would say for sure, was digging and poking at looking at different kinds of methods, and then in comes this little program, and I take the test, and what appears is a name. At the time, we were using names, because I don't think the program could handle pictures. So, a

name like Tyrone or Jamal would show up, or a name like Tom or Joshua would show up, and I would have to classify them. Use the E key or the left side of the keyboard to respond to Black names, and the right side of the key, use the I key on your keyboard to respond to white names. Super easy. If I see a name like Josh or whatever, Steve, or any one of your standard Anglo-Saxon, Judeo-Christian names, I'm supposed to hit the one button, and if I see a name that comes from a subset of Black people, then to distinguish it from the names that were clearly different, then I press the left key. How hard can that be? Super easy.

And then I do the same thing now with words, and now the same key that I'd use, remember, E was the key for Black names. Now, whenever a bad word pops up, a word like poison or devil, I'm supposed to hit that same key, and then when it's a good word, like love or peace, I have to press the right key. Easy. I can do that.

Now, the program says just put them together. We'll show you one of four things. A name will pop up that will be clearly either Black or white, and then a word will pop up that will be clearly good or bad. Whenever you see a Black face or a bad word, press the left key. You've already practiced these. And whenever you see a white name or a good word, press the right key. I do that and my fingers are flying on the keyboard and I'm done. I'm done in like 500 milliseconds and it pops up and no response, just told me that my average time was like 535 milliseconds or something, so that felt right.

And then the other version comes along. Now, you'll switch. Now when it's a Black name or a good word, press the E key. When it is a white name or a bad word, press the I key. Okay. That seems like it should be... I mean, I remember this very clearly. I remember starting that task knowing, knowing that I would whip through it in the same time.

Andy Luttrell:

And where were you when you took this? You said you got this from Tony?

Mahzarin Banaji:

Yeah, I remember it well. I lived in a little apartment in New Haven, in Connecticut, and the dining table was my office, and I had put a desktop machine on the dining table. And I remember taking it there. And I mean, all of a sudden I can feel that my heart rate is going up, and my hands have just become completely clammy, and so I pause, and I think, "Okay, just clear your head. Go back." And I start, and I can't. I can't do it. Of course, I can do it. But much slower and with many more errors if I try to speed up. In those days we didn't... It was a prototype, so all it did was spat out some number in the 800s or something like that, and so my first thought is something is wrong with this test. It can't be this.

So, of course I went where every other person has gone. It's the order.

Andy Luttrell:

The number of times I've had to tell students that that's not an issue.

Mahzarin Banaji:

Yeah. And then we did work further. In the old days, we didn't have it all taken care of, but now we know that we're wiping out all order effects if there are any, and there are many tests where the opposite order actually produces a bigger effect, so we know it's not any one order. So, I can poke around in the program and I can take it any other way, and no difference really. So, I know that I just... I think the reaction was... I call it the most transformative day of my life, and it is true that no other single experience has ever had this kind of impact on me. But I knew that this was incredibly important.

So, I asked my husband to come and I just say, "Take this test." Our families had been used to being guinea pigs in these tests, because we're among that very lucky group of scientists where we can be subjects in our own studies. If you're a physicist, you cannot become an atom, and if you're most kinds of psychologists who know what the study is about, it won't work, so we're like people who study optical illusions or something like that, because we too will show the same effect that our subjects do. So, I strapped him into the chair and said, "Go."

And he finished it and he said, "Don't ever share this with anybody in the world. Ever." And I understood what he meant. I felt like yeah, maybe. Maybe that's what we have to do. We have to just hide this. And then of course, two seconds later, I knew. We should put this thing on a website.

Andy Luttrell:

Well, what was the concern?

Mahzarin Banaji:

You know, I think if... This is where I think it's so hard that once you're in quantum physics, it's very hard to understand what the world would have been with just Newtonian physics. But when you ask that question, it's a legitimate question, and I love it, because to you, it seems like, "What was the concern?" And you didn't do that kind of thing. You didn't do science on race. You didn't do that kind of thing. You stayed away, very far from those kinds of things. There was science and then there was politics, and this was about politics it seemed, and what are we doing with it in the science? So, there was that, but much more than that was the feeling we've discovered something that is important/scary, and then what to do with it, and so I think my husband's response was a very caring one. You know? "My wife's life is about to change if this gets out."

And I didn't feel that way. I didn't feel like we shouldn't share it. I knew immediately that this was extremely important. I think I took it in the summer, and that year, '94, we went to SESP and we didn't have laptops in those days, so we took a bunch of playing cards, and we made an IAT on playing cards. So, we would take each card and on it we would put a name, like Tyrone, or Josh, and then words that were good or bad, so the cards would be divided into four categories. They would have on them good or bad words, and then what were distinctively Black names or white names. And we would shuffle the deck. We would bring people into our hotel room and we would say, "Okay, sort. We're gonna start a stopwatch. Sort." Basically, we'd make them do the IAT.

And when it would come to that second moment, I mean most people just would throw the deck of cards on the bed and turn away, because they could feel that they can't do it. And then we would laugh, and we would say, "See, you see this?" And then they would say, "It's not what you think,"

and they would leave the room. And then there were some people who were like, “I’ve always known this.” I’ve always... I’ve been so grateful for those people. “Of course, I’ve known this. This is not a shock to me. Yeah, I’m very glad you’re showing it, but this doesn’t surprise me.” So, yeah.

Andy Luttrell:

So, what is it about that test, that sequence of things, just to sort of look under the hood a little bit. Why is that revealing of anything?

Mahzarin Banaji:

Why is what revealing of anything?

Andy Luttrell:

Really, I’m just asking to explain the logic of the test and what it’s really measuring.

Mahzarin Banaji:

Oh! Oh. The IAT. Yes. All right, the IAT is doing something really simple. There is a fact of mental life that when two things come to be paired in your experience over and over again, that you will respond more rapidly to them. Salt and pepper go together, so when I say salt and then automatically I’m gonna think pepper. And so, for about 50 years before we come on the scene, cognitive psychologists have written books on response latencies as a measure of associative strength, so we are not even new there.

You know, we’re really totally unimaginative, uncreative people, who took existing stuff and put it together in a particular way, and when none of those individual strands had produced any reaction at all, putting them together did. So, people did say, “Well, what is it then?” So, we would say it’s the strength of association. If bad and Black are so much easier for you to put together than white and bad, that means that like salt and pepper, or king and queen, or sunny day and beach, or whatever it might be, this is just telling us what has been paired for you in your life. That’s all it is. It’s not saying you will behave this way or that way, but we didn’t even go there in those days. We just said, “It’s the strength of association.”

Andy Luttrell:

So, just to maybe shift gears a little bit, we’ve talked a lot about what implicit bias is, how you’ve measured it, whether it’s something people are aware of or not, and I thought maybe for the rest of the time, what we could do is talk about how changeable you think these implicit biases are. And in doing so, continue to keep in mind what the differences are between implicit and explicit biases.

Mahzarin Banaji:

Sure. So, this too has a very straightforward and very interesting story. When you have the experience that Tony did, being the first person to take it, and me being close to that, it’s very... It should be very easy for you all to understand why we thought that it was not gonna be changeable. Because when you know what you believe and then how the data are coming out, if that’s happening, what other than some silly thing, like I’m just gonna practice something and then show that I can beat the test... It felt just to me like if this is the case, the world is gonna have to

know this and if there are to be policy changes, they're gonna have to come in the form of assert them from the top. You're not gonna get people to change. So, you're gonna have to do something else, because people cannot change. And this is implicit. This is ground into us in such a way that it's not gonna change in my lifetime.

So, that was my incorrect assumption. So, students would come to me and say, "We'd like to do a study to see if we can show people positive exemplars, this and that, and shift the IAT." And I would say, "No, you're not gonna do that on my watch. I'm not gonna waste government money on doing something that is totally useless. You can do that when you leave my lab." And sure enough, three of them escaped. Curtis Hardin, Irene Blair and Buju Dasgupta, who went actually to do her postdoc with Tony. And each of them, within a six month period, came up with data that in one way or another showed that you can change them. And you know the studies. In Buju's case, you show exemplars of positive Black individuals, like Martin Luther King, Jr., athletes like Michael Jordan in those days, et cetera, and you can diminish the magnitude of the IAT race bias.

Curt Hardin did a study in which he just a Black or white experimenter come before a class and do a paper pencil version of the test. And he found that with the Black experimenter, the race bias in white American students, but not Asians, was lower than in the white experimenter condition. And Irene Blair told college seniors roughly to just spend five minutes in their heads imagining their first job after college, and that their boss was a woman, and they spent that time doing it, and then she gave them a gender stereotype test of male, female, strong, weak, and showed that if they had imagined a female boss, they showed a less strong gender stereotype.

These three papers all appeared together in an issue of JPSP, and in retrospect, I was just shocked that I hadn't seen it, because there was a graduate student who had come to me and said... You know, we were all, in those days we would take IATs all the time. We'd make them up and send them to each other, and we'd be always taking them. And one of my students said, "I'm gonna..." He was so outraged by his race bias that he would come in every day in the morning to the lab, he'd take a race test, and he would write down his score on a board, and he would track it over time. Everybody knew not to erase that part of the board. And he came up to my office one day and said, "Hey, something weird happened. Every day I've been showing a race bias. It varies a little here and there, but it's been pretty much the same. But today, I showed no race bias, and so I thought to myself, what did I do before I came in." And he was a runner, and he said, "I was watching Olympic track and field.

And I said, "Wow, isn't that interesting?" But I totally let it go, because my belief was that this cannot change, and that this was just a little anecdote, but I didn't use it as the beginning of a hypothesis that maybe, maybe something is telling us that it can change. Now, young people came in after those studies had been done and they went to an extreme that I still have trouble understanding. They got all upset when the bias that had changed did not remain changed. They were like, "God, the Buju effect is gone after a week." What did you think was gonna happen? There's a thing called a world out there. Imagine, your brain is adaptive, it's changing when it's being given some signal that something that it had thought was true for a lifetime may not be true. It's so amazing that it can change. That it snaps back is what an adaptive system does. Go think about this.

But oh no, they were all huffy about how this was all useless, because nothing. Because they're thinking politically. They're thinking, "Oh, this means the world will never change." But you can't be that way. For you, it's just the science and whatever it's telling you that should be important, and if it's telling you it can't change after time has passed, you should use that data to revise your theory about how the mind works, not to say, "Oh God, the change is not lasting." So, when we would write about this, we would call these changes, we would say so far we know that they're elastic.

And then a lot of people did a lot, a lot, a lot, a lot of studies, but they were all incredibly weak of this kind. Showed them five exemplars of some counterstereotype and then wring your hands when the bias pops back to where it was. And I to this day am irritated at this group of people who wrote so many papers saying implicit bias, no change over time, after having done the weakest studies you could imagine and making this claim.

Andy Luttrell:

And your frustration was that it was a criticism that they're saying.

Mahzarin Banaji:

Yeah. Yeah, it was a criticism. It shouldn't have been... It should have at least been just, "This is what the data are telling us, that this does not last very long," which is how I would have said it. And I would have said expectedly so. Can you imagine it actually sticking? You don't want a brain where every little thing you do is going to remain forever. One instance should not matter to a smart, intelligent system. So, having said all that, I can now tell you that the best data that we have that is plastic comes from Tessa, and we have not done sort of the classic experiment of bringing people into the lab every day and then showing that after six months or a year of this kind of training, that no matter when you give them the test, the data that have moved have now stuck in this new location. Nobody's done that study. The resources it would take to do it are serious, and that's why people did these easy studies that were actually weak studies.

But then, of course, we have two things going for us. We have data being collected in a way that no other data set has been collected. There is no survey data, no any data where continuously, morning through night, every second, every minute, data are being collected. To me, that's the most striking thing about the IAT website, that it is without stop collecting data.

Andy Luttrell:

It's almost like a weather center, right? That constantly monitors the weather.

Mahzarin Banaji:

It is like... I will use this, and I will give you credit for it, but it is a great metaphor. It is like a weather center. And we believe weather data more because it is doing all this averaging, minute by minute, of what's happening. But of course, we didn't have the methods to analyze data like that, and of course there are huge alternative possibilities if you see change. Cohort differences. More younger people coming now. The one thing we learned very quickly, in the beginning people would tell us, "Well, your people who come to the website, to Implicit.Yale or Implicit.Harvard are coming to it because they're interested in bias and they tend to be," our data showed they were

more liberal than the population and so on. But very quickly, that changed. Even in the first six months. We would look at people coming to the website. Like eagles, we'd be watching.

So, what this did for us is it made the population coming to it quite representative of the country. You know, now the federal government will send the entire Justice Department or all of ICE to take these tests. So, you get people across the spectrum, so we have now data that are extremely representative. Huge, compared to any other data set. Millions of people. Enough that we can do analyses by zip code, and it's continuous. But still, I didn't know how the hell we would analyze these data. And then we had hired at Harvard a wonderful statistician, and we asked him, and he said, "Yeah, you should look at how economists measure. Their data are not necessarily as good, but they have huge data and they're fixed in this way, and they're over time." And those are time series models, and we could tell that if applied properly, this would actually be the right way to analyze these data.

And those data showed that on sexuality, there's a huge, huge shift. 50% drop in explicit bias, but 33% drop in implicit bias. In 2005, anti-gay bias was among the strongest biases, and now it's suddenly among the weakest of biases implicitly. No other group has changed quite like that, but again, just as I have my irritation at the people who say, "Why didn't the change last?" In this case, I'm gonna make a very big deal about this change, because others are not showing the same pattern. Right? So, the race test shows change to the tune of about 15% over the same time period, a decade, along with skin tone, but also no change on age, no change on disability, no change... In fact, some evidence that it's gotten worse on body weight bias.

So, you've got now four different patterns using the same method of analysis. Rapid change in implicit attitudes, modest change in implicit attitudes, no change in implicit attitudes, and a reverse change, reverse in the sense that the bias has become stronger. So, this is very interesting, because it can't be that our method is just picking up some pattern and replicating it. If all of the effects looked the same, we would be suspicious. But the excitement is that something did change hugely, and now the question is why. And we've been on that path for almost a year and more now, trying to identify all the pieces of sexuality that differentiate it from other biases. I'll mention two quickly. One is of course segregation. Gay people exist everywhere, amongst poor and rich, West Coast and East Coast and Midwest. They can be your neighbor, in your class, your cousin, your sibling, et cetera. That's simply not true for race. We are deeply segregated, and especially on the three things that most matter for segregation: where you live, where you go to school, and where you work and what you do there. That's one feature.

But we know that contact doesn't always produce positive attitudes, and here I would say that what sexuality had going for it was that it was socioeconomically also varied. So, there were rich people and there were people in Hollywood who made movies about gay characters who were cooler, and smarter, and better in every way, and I think over time, it didn't just change some minds. What's remarkable about the sexuality test is that everybody changes. Elderly people change and young people change. Gay people have changed, and straight people have changed. People in all the regions have changed. Very rarely does this happen, that the country as a whole is shifting. And I think that there is one or two counterexamples, and it is not... What I just said is true, but it is true that while everybody is changing and going along exactly that little path of moving towards zero, towards neutrality, there are two groups that are changing faster, and those are people who claim

to be liberal and especially people who are younger. But that's not to say that other people are not changing. Everybody is changing, and I think that's worth us knowing, because I'd like to ask the country why that is not happening for every other group.

Andy Luttrell:

Can I ask just to clarify those data, you said it was both implicit and explicit responses changing, so I'm curious, what is it that is unique about that, that both of them are changing? So, part of a skeptical approach would be to say, "Well, if all we needed to do was ask people explicitly and we could track that change, that would be easier." Why is it interesting that the implicit seems to be doing the same thing?

Mahzarin Banaji:

Yeah, because we had seen lots of change in explicit. Remember, from the 1933 Katz and Braly data, we had seen explicit fall all the time, enough that you couldn't ask people the same old-fashioned racism questions. They'd become so meaningless that everybody was topping out, getting a high score of 100 or whatever was the highest score. Everybody was looking completely not racist. And yet any look at the real world and you knew that was not true, so it's kind of a little bit Sherlock Holmesian, when every possibility has been exhausted, what remains before you has to be the truth. And what remains before us is that something implicitly was happening.

So, not all things are changing equally for explicit, implicit. I can share those data with you. I just put the sexuality ones out there. I think for race, which was already kind of more modest. People were willing to say they were anti-gay but not anti-Black in 2005. So, the changes in explicit are not always equal to changes on the implicit, and we have a section in that paper on what the patterns of change are for explicit and implicit, and sometimes the explicit will stay the same, because they're already pretty close to zero. But the implicit will continue to change, because there's a lot more room for the implicit to move.

But they don't ever go in opposite directions. I don't think you I've ever seen data where the explicit became more negative and the implicit became more positive. Which again, was the same. They are linked, because the same experiences are feeding both these systems. It's just that it's not the same because of the way our brain stores different types of information, retrieves it, uses it, and how we behave, that these are different enough that these contexts make enough of a difference that implicit and explicit are different from each other. They're just not opposed to each other I would say in most cases.

So, take something like political attitudes, where the correlation between implicit and explicit is like 0.8. I would say there, don't bother with the IAT. Just ask people, "Who are you gonna vote for?" Because what you learn from the IAT is identical. But you become a little more complicated and you'll see why the implicit measure is incredibly interesting. We have a very large group of people in this country who say that they're independent, but independents rarely show no bias. They do veer in one or the other direction. Very few are zero. So, there's something interesting now. You're saying you're an independent, that's your conscious belief, but we can show you you're not. You are a little bit this way or that way, and if the right people did the tests and they tracked what the independent said, and then what the independents showed on the IAT, I wouldn't be surprised if the IAT then predicted how they voted.

So, the one thing I want to say, because I think it's also misunderstood, and people who wrote in this field, and I would put myself in that category, we had a simple understanding early on that maybe there are two worlds, the world of conscious attitudes that affects my conscious behaviors, and then there are these implicit attitudes, and those affect my more implicit behaviors, like eye gaze, or seating distance, or something like that. And a lot of studies were done on that, and I think that was a waste, because it's not quite so simple. It's actually much more interesting. My understanding today, the way I would state what I think happens is that the implicit attitude affects eventually your conscious attitude, and it's your conscious attitude then that determines what you do. But the conscious attitude becomes informed by the implicit.

If I make less eye contact with my Black students than my white students, if I'm a little more suspicious of my Black students performing very well on a test, each of those little implicit feelings will bubble up in some way and turn my conscious actions into the actions they are, where I will say... I'll say, "You know what? I just don't think that it's worth calling on Jamal as much as it is to call on Harry," or something like that. So, this is a slightly new idea, and I don't think we've actually demonstrated this, but my sense is the way to think about this is that yeah, they may have their own little spheres of application in the way William James talked about, but I think we're gonna be able to show that the implicit feeds into what becomes the consciously held belief, that it feeds it and shapes it in some ways. And a lot is going on. The implicit stuff is shaping it in one direction, and then our conscious values and everything else, that is part of the true conscious system is also shaping them. And that's where the mix happens in some way that we will in the future learn a lot more about.

Andy Luttrell:

And an implication of that too is that if a road to changing those conscious actions is to try to manipulate those implicit biases, so as a segue into kind of the last thing that we were gonna talk about, over the last several months, we've seen lots of renewed discussions about racism and police brutality, which have ignited, as you've said, sort of strong opinions on either side about what implicit bias training can do and is appropriate to use as like the answer to these sorts of issues. And so, I'm just curious to get your take on what those what you characterize as extreme positions on the issue are, and what that middle ground for you actually looks like.

Mahzarin Banaji:

So, it's actually very similar to my description of earlier of mostly young people's response to when the IAT change snapped back. They were puzzled, they were sad, they were upset, and here, when you do something called go give a lecture on implicit bias to some military group, or a corporation, or a governmental agency, and it doesn't look very different a year or two later, I think one group of people have a very similar response. It's so sad. Implicit bias training didn't even make any difference.

And my answer to them is imagine that you had spoken to them not about implicit bias, but about something like body and health, and you said, "Look, this is how our bodies are healthy, every day eating the right foods, every day exercising, reducing stress, these are sort of the three things we know are important." At the end of a three-hour lecture on body and how it metabolizes sugar and fat, would your audience have lost any weight? And everybody laughs, because of course they

shouldn't, so I say, "Why do you think it would be any different with this? You are telling them about this to explain to them what's actually at stake here, what might be going on, why they think that they're unbiased when in fact they aren't. And then they or you have to figure out a way to get them on the equivalent of the treadmill. To eat salad more." Whatever the equivalent is, it's in the information coming to them, right? So, what food is for the physical body, data and information are for the mind. So, now, I've changed. I don't just say, "I'm gonna just watch a movie on X," or I don't just go to a particular website. I actually think about what I want my brain to learn and what I don't want for it to learn.

And I feel that my own changes on the IAT are almost a direct reflection of whether I've worked to change the input into my mind. So, the first thing is of course you shouldn't expect that to happen. Again, it's the same thing. It's a drop. And it's not even a drop in change, it's just a drop in alerting somebody, so why should they suddenly start changing their patterns of who they hire, and who they promote, and so on?

There is another group that says since nothing changes with implicit bias changing, we shouldn't be doing it. They'll say, "Let's not do this. Let's just change the world." So, they're saying just go ahead and hire people of a certain kind. Have quotas. And I would say then you don't know anything about human behavior, because when an organization or a government comes in from the top and says the data are clear, there is bias, you're not gonna be able to control your bias, so we're just gonna shift the defaults in how we hire, what we do. If you do that, then your audience of individual humans in that system are going to respond like a three-year-old being told to eat peas and carrots. They will spit them in your face, and with good reason. You are removing them from decision making. You're telling them that they're no good and that you're gonna do something above their heads to make change in your society.

So, I have argued that there are three levels of change, that the first and the one that psychologists know the most about is change at the individual level, an individual mind changing based on inputs, framing, all the things that you know much better than I that go into attitude change. But if that happens and there is nothing at the next level, which I'll call the mezzo level, if nothing happens at that middle level, which I would argue is the level of the institution, your university, your community, your children's school, your town, whatever is that middle level, if those people don't engage at the same time as you are engaging as an individual on this change project, nothing will happen. Because you can change your mind, but if your hiring practices still use interviews with all their biases in the way they hire, nothing is gonna change. It will be like teaching people and then not putting them on a treadmill for body.

Andy Luttrell:

This is kind of like the data that you have, the zip code kind of IAT data, to sort of characterize a community.

Mahzarin Banaji:

Yeah, exactly. And because we are changing that way, because you know, places with higher anti-Black bias are the zip codes where there is more lethal use of force by police or whatever the other data that my economist colleagues have collected using the IAT, showing less upward mobility, et cetera, et cetera. But it's more than that, I think. It's that only institutions have the power to change

processes and policies. You and I could be going into work every day, and if our university tells us that we must look at the GRE scores of a candidate first because that's what they put in the top of the file, and then everything else, we are gonna continue along a certain path. But only the institution can decide to say, "You know what? GRE scores used to be predictive when the world was pretty homogenous. Given different peoples experiences and so on, they've become less and less predictive, blah, blah, blah, we're gonna change that. You won't see the GRE unless you want to. It won't be right in your face, but if you want to look at it, you can." Okay, now that's a change that I as an individual cannot make, and I'm gonna argue that unless you've got the organization doing its bit based on the same data about implicit bias, it will not change. But you put these two together and I think you would see fast change. Massive change.

And of course, these are the only things that are under our control. The third is what we are going through. Something happens at the level of a society. A society sits quietly and just sucks it up, and then there is a little break, and a little revolution happens, and when that lines up with individual change and institutional change, I think you cannot be in a better moment. Because as you are doing your individual change and as your organizations are changing at the level of the institution, and your society is changing the names of a sports team or removing a statue and so on, then all these three are moving in a single direction, and I think that's what happened for sexuality attitudes. That individual change based on cognitive dissonance got resolved in the direction of I'll change my attitude to gay people, institutions coming out, greater acceptance, et cetera, saying that people who cohabited could get some... And then of course, society. Change in gay rights, gay marriage, and I think we've seen that we can make change if we do it at all three levels, or at least I have a model for why I think anti-gay bias can be our model, because we have made it work, and not many societies at many times can point to evidence of the kind we have and say, "Look, it happened here. Why? Why did it happen and why can't it happen elsewhere?"

And I think we're in that moment, and for that reason, my email box, inbox, is filled with calls from people as diverse as garbage management and ballet companies calling to say, "What can we do?"

Andy Luttrell:

To bring it to what it means for implicit bias, it kind of sounds like you're saying that changing those individual biases is necessary, but not sufficient.

Mahzarin Banaji:

Yes.

Andy Luttrell:

Okay.

Mahzarin Banaji:

That's pretty much what I say. It's necessary, because I believe, I guess I'm just sort of a person who believes in democratic free speech, openness, I don't like systems that impose things on individuals, I don't think they work, that you cannot take away the free will of the people, but you can nudge the people with your own actions and tell them that they can choose freely, but that in your implicit bias training, you can show them that what you're showing them is not only real and

not made up, but also that it might be in their interest, that they may not want to wait until people come for them with pitchforks, that the banks may want to change how they do their business. And we have changed. The same people, our ancestors did believe that slavery was a way of life, and we stopped believing that, so how can we say that change won't occur? Change has occurred. The question is is it in the right direction and is it moving as fast as it ought to?

Andy Luttrell:

Well, that's a very hopeful note to end on. Mahzarin, thank you so much for taking the time to talk about your work, and where it came from, and where it's going in the future.

Mahzarin Banaji:

Thank you, and this was fun for me to do, and I noticed that unlike many other interviews, partly because you're a social psychologist who knows what I know and shares all of that, I felt like I was able to tell sort of the story behind the story in many cases that I have never really talked about. I didn't say something, but the website is an important moment. 1998. Four years of quiet just doing inside work in a couple of labs, Tony's and mine, and then we decided to go public.

Andy Luttrell:

Was that before the paper came out?

Mahzarin Banaji:

It's the same year, so the paper comes out and the website. I think the paper may have come out earlier in the year, and September 29th, 1998 is the opening of the website, and I remember telling Tony and Brian Nosek that wouldn't it be amazing? In a year, we could have 500 people who've completed a test and we could write a paper with an N of 500. And just report, and no longer will they be able to say, "That's just because they were Yale students, or UW students, or whatever they might say. This will be people." And in the first month, we had 45,000 completed tests. No advertising.

Andy Luttrell:

What was the moment that allowed so many people to know that that existed? Like you said there's no advertising, but where did people hear about it?

Mahzarin Banaji:

I think there was a small press release and a little article might have been written for Yale, some Yale magazine. I remember that we did put just two paragraphs, you know, psychologists have been studying implicit, blah, blah, blah, for a while. We just put that out, but I don't think people went to it because they read that. But they... It was word of mouth.

I'll tell you a very funny story. About a year after the website had opened, I was on a plane coming back from a conference, and sitting next to me on a small plane was a man, and I was doing my work, and he must have looked at something and he said, "Are you a psychologist?" And you know as well as I that when we're asked that question, we want to become miniscule and just hide, because we know what the next question will be. My girlfriend dumped me, and blah, blah, blah. So, I sort of hesitantly said, "Yes." And he put two hands out before him. He said, "Do you notice Black, white," and he pressed his little hands in the air. "Do you know this Black, white, good,

bad test?” And I said, I thought, “I guess I’m on candid camera. They’re gonna tell me in a minute.” And so, I said, “Yes, I knew that test.” And I said to him, “So, who are you?” Assuming he’s some kind of college student whose teacher made him take it, and he said, “I’m a pig farmer.” I said, “What are you doing?” He said, “I’m going to Boston because in Western Mass there is a pig festival that all the pig farmers are going to.”

And I said, “Who told you about it?” And he said, “Oh, another pig farmer. He sent it to us on email and said we should try this.” And he didn’t seem bothered. He goes, “It was kind of interesting.” And that’s when I knew this is not in our hands anymore.

Andy Luttrell:

Yeah, it’s in the world now.

Mahzarin Banaji:

Ah, yes. For good or for not, but I’m hoping for good.

Andy Luttrell:

Well, it was super great to talk about this stuff. It was really neat to hear the stories behind all the different aspects of where the notion of implicit bias come from, so Mahzarin, I just wanted to say thank you so much for coming on the show.

Mahzarin Banaji:

Yeah, and thank you.

Andy Luttrell:

That’ll do it for this episode of Opinion Science. Big thanks to Mahzarin Banaji for taking the time to be on the podcast. We talked in general terms about the results of her research over the years, but if you’d like to know more about the specifics, check out her book, *Blindspot*, or check out some of the links in the show notes. And hey, if you enjoyed this episode, please head over to Apple Podcasts and leave a kind review. Rating and reviewing the show helps get the podcast into more ears, so thanks for your help with that. For more about Opinion Science and to get a transcript of this episode, head on over to OpinionSciencePodcast.com or follow us on Facebook or Twitter @OpinionSciPod. Hope you’re all doing well, and I’ll see you next time for more Opinion Science. Bye-bye!