THE BROOKINGS INSTITUTION

WEBINAR

RACE, ARTIFICIAL INTELLIGENCE, AND SYSTEMIC INEQUALITIES

Washington, D.C.

Friday, June 19, 2020

PARTICIPANTS:

Moderator:

NICOL TURNER LEE
Fellow, Governance Studies
The Brookings Institution

Panelists:

FAY COBB PAYTON
Professor of Information Technology and Business Analytics
North Carolina State University

RASHAWN RAY
David M. Rubenstein Fellow, Governance Studies
The Brookings Institution

DARIELY RODRIGUEZ
Director, Economic Justice Project
Lawyers' Committee for Civil Rights Under Law

* * * * *

P R O C E E D I N G S

DR. TURNER LEE: Good afternoon, everybody. I am Dr. Nicol Turner Lee. I'm a senior fellow at the Center for Technology and Innovation at the Brookings Institution. I'm glad that you all could join us, particularly on this day of Juneteenth.

Now, before I get started, I know that there are some of you that do not know what Juneteenth is, so I'm going to make sure you know before we start the conversation. And if you want more, my colleague Rashawn Ray is up on the Brookings Instagram giving a definition around it.

But Juneteenth, on this day in 1865 is really the day when the last of the slaves in the United States were actually freed in the Confederate South. And that's really significant because Lincoln, with the Emancipation Proclamation, had declared us free two years prior, but there were a group of people in Texas that did not know. And when it was finally announced to them that slavery was over in the U.S. for people from the African diaspora, it was a day of celebration.

And some of you may be wondering why are we celebrating Juneteenth on this day or why are we particularly excited about it as we have gone through really this very difficult time in race relations in our country and civil unrest? Well, I'm going to tell you why on my perspective and then you can disagree if you don't like what I say. Just tweet it. But I'm going to tell you the reason we're so excited is because for one of the few times in history that African Americans get to write the good parts about this story. And that narrative is so significant because part of why racism exists today is because the story is not necessarily always based on everybody we're talking about.

So, with that, I'm actually excited that I didn't know that I was actually scheduling this particular seminar on Juneteenth, but I'm glad I did. I also didn't know that I was going to be scheduling time of such unrest and such heightened discussion about race, but I'm glad I did.

And so, at Brookings, I work on issues related to digital divide and digital inclusion. Shameless plug, I have a book coming out next year on the digital divide. But I also spend a lot of time looking at what's called algorithmic bias. And I think you're here because one of the issues that I really wanted to talk about for a long time is the issue of race, racism and systemic inequality, and AI.

And it's no secret that in every conversation that many people are on, particularly those

that are highly engaged in this space, that we always come back to the potential discrimination that can happen due to the next innovation in technology. How does that happen that we're talking about either the marginalization of groups when it comes to employment AI or facial recognition technologies creating further bias in the overcriminalization of certain groups or in healthcare where we're seeing bias actually come out of algorithms that just most recently kicked out Black patients that were deserving of a spot in this program because of chronic disease?

So, I'm joined today by three people that I absolutely just respect their work. And I'm going to introduce them as soon as I find my paper, which dropped on the floor, in the order not necessarily in which they will speak, but just to give you a sense of who is with us today.

Today I'm joined by Dr. Fay Cobb Payton, who's a full professor of information technology systems at North Carolina State University. Rayshawn Ray, who is my colleague and also a David Rubenstein fellow in Governance Studies here at Brookings. And Dariely Rodriguez, who is the director of the Economic Justice Project for the Lawyers' Committee for Civil Rights Under Law.

Thank you, everybody, for being here.

Fay, I'm going to jump right into this because our time is short. In recent weeks, as I've just mentioned, we've been so sensitized around the discussion of race.

Oh, before I forget it, please send questions to events@brookings.edu or @BrookingsGov or #AIBias. Events@brookings.edu or @BrookingsGov or #AIBias.

So, Fay, let me go back to you. Race has been a heightened and sensitized conversation, right? And I'll need to make sure everybody unmutes themselves. Particularly as we looked at police brutality and other forms of race have actually surfaced, right, among corporate America, we're actually seeing it, you know, in terms of the systemic inequalities that arise around visual equity.

These are not new issues. And as James Baldwin would say, in all of the beauty of the United States, it's the terrible conscience in history of this aspect of our society that's been most daunting.

I'd like for you, as a person who does computing and engineering, to help us unpack why race and racism in some cases factors into artificial intelligence systems.

DR. PAYTON: Thank you, Nicol, and thanks for having me on the panel today. So, let

me start out with this piece of a narrative.

Unpacking race and racism is the question and why is that critical? The idea that it's critical is because ofttimes technology is designed to be for the betterment of society. What I say is ignoring factors like race and racism does not necessarily eliminate those issues. Those issues don't go away. And, in fact, often heighten bias, whether explicit or implicit, that can be built into an algorithmic function. That means that oftentimes that science may not necessarily be asking the right questions or addressing the best issues or the most salient issues.

So, if you take a look, science drives and data drives most of what happens in AI. And that means that it's driving what's recommended, who's influencing, and the policies that are associated with lived experiences, particularly of Black and Brown people. That is of no -- that's obvious, right? So, the exclusion of the populations in the algorithmic design process from the beginning holds implications from the science in the end.

So, take for instance in 2016, there's data that shows that -- and I have this written down -- that 9 percent of science -- bachelors of science degrees that were awarded were to African Americans and 13.5 percent were to Latinos, but only 5 percent of the Ph.D.'s in science and engineering were to women who were Black in that same year. And that was for women who were Black, Hispanic, Latina, Native American. And for men in that same group it was only 3.8 percent.

So, when we're talking about this idea of representation and training the algorithm to understand the lived experience, not reading about the lived experience, only seeks to complicate these matters and why race and racism is really important.

So, lastly, to give you an example, particularly in healthcare, did a study with a number of colleagues from engineering and health policy. We took a look at focusing on mental health and HIV. What we found out was that these conditions are actually comorbid. They exist. They cling together. They are not in isolation.

So, looking at mental health and disorders it said that there are cost of care and total charges that obviously resulted from the model, but when you don't factor in race, you get a totally different picture of who actually gets services and how many services they will get. So, it's not just getting

the services. It's the quality of the services that are to be given.

So, I say, and I'll end with this and pass it on to my colleagues, is that in the end, we see that omission of race and racial lives' experiences continue to challenge. And so I want to bring this home with a second healthcare example.

In 2018, Linda Villarosa wrote a piece in *The New York Times Magazine* that talked about why are Black women and babies in the U.S. dying at the rate of more than two times the rate of White mothers and babies? And who gets to be healthy then is a topic of Time magazine just last year. And so, when we're talking about race and racism, and in healthcare in particular, we want to talk about precision medicine and personalized medicine, we can't do that unless we're given recognition of race and racialized experiences in the algorithmic design process.

DR. TURNER LEE: Yeah, and, I mean, Fay, on that point, we've seen it just a year ago, right, with the racial bias and that algorithm that was structured to target people for a particular program around chronic disease.

DR. PAYTON: Exactly.

DR. TURNER LEE: And because the variable was the cost of care and, historically, African Americans pay less into their care --

DR. PAYTON: Right.

DR. TURNER LEE: -- they were kicked out of the program --

DR. PAYTON: The program.

DR. TURNER LEE: -- despite having higher rates of chronic conditions.

DR. PAYTON: Absolutely.

DR. TURNER LEE: Yeah.

DR. PAYTON: Absolutely, yeah.

DR. TURNER LEE: So, I think that matters. Uh-huh?

DR. PAYTON: It does matter. That was actually my next example for my next question, right? But that's okay. So, you preempted me, right? And so, the idea that we're using proxies, right?

DR. TURNER LEE: Yeah.

DR. PAYTON: And the idea of the proxy means what kind of care will you receive, if at all. Right? And if you're kicked out of the algorithm, it's interesting how some algorithms kick us out and some include us. Right? And when there's inclusion, ofttimes it's not necessarily that the outcomes will be favorable to people that are Black and Brown.

DR. TURNER LEE: Thank you for that, Fay.

And so, Rashawn, and for those of you that are wondering, this was on purpose an interdisciplinary panel or some of the research that I'm actually testing right now, the extent to which you can have more robust conversations when you bring people from different disciplines. So, Rashawn, I want you to unmute and I want you to actually go into as a sociologist, you know, sort of unpack what you and I were sort of raised on in terms of our training, which is race is a social construction, one that is primarily defined and narrated by where you stand in the power spectrum.

When you think about algorithmic design and you think about algorithmic models, and I know you do that both at Brookings as well as at your center at the University of Maryland, how do we look at race in this social construction and what Fay is sort of starting in this conversation, the design, the implementation, and execution of algorithmic models?

MR. RAY: Yeah. I mean, I think what Fay said was excellent because when we think about algorithms, they come from somewhere. I think one of the assumptions is there's some sort of way that math or statistics or computers are some kind of way free from bias, that they are objective. But they are created by people. And as Fay was mentioning, what goes into the algorithm informs what we get out of it.

So, the disparities we've been seeing in healthcare, I mean, when you talk about COVID-19, Black people are 6 times more likely than Whites to be turned away from testing once they get to the hospital. That's because hospitals are using an algorithm either that someone is making that decision for or they're entering people's information into a computer. The bottom line is that Black people are more likely to be sent home to die.

And I think it speaks, when we talk about technology, that there is a huge gap in trust in how we think about facial recognition, for example. So, Black people are about 33 percent less likely than

Whites to trust facial recognition. Well, why would that be? That's because studies have shown that Black people, Black women in particular, that Black women have been misclassified about a third of the time. I mean, Amazon just put a halt on their facial recognition program, Rekognition, that has been incorrectly labeling, say, professional athletes in Boston and interpreting them as being criminals.

And now think about that in this moment, about the way that when we think about the social construction of race and also racism is in the application of that. That right now all of these technologies are being used around the country as people are protesting in the streets. And I think we really have to be clear that as we think about doing this on Juneteenth that there are people right now, whether that be in D.C. at Black Lives Matter Plaza, all the way across the country and around the world, who are aiming to raise up and liberate Black people who continue to be marginalized in the United States and around the world.

And with that being said, we know that law enforcement is using advanced forms of technology to surveil people. And we know that this surveillance is oftentimes applied inequitably.

So, if we go back to what happened at the University of North Carolina with the Silent Sam statue, you had two groups of people. You had Confederate people who were trying to save the statue and then you had people who were going against the statue for saying that it's racist. And interestingly, the police only used geofencing on one group of those people: that was the people who were actually arguing for the removal of the statue. So, we see the way that technology is discriminately applied based on who is actually engaging in a particular act.

And so, when I think about the social construction of it, part of the fundamental problem, particularly when I talk to computer scientists, overwhelmingly when I talk to statisticians, is that they really approach algorithms and technology in a color-blind way. They say, oh, okay, if we don't talk about it, if we don't mention it, then we ensure that race is not part of it. Like that couldn't be farther from the truth.

In order to fundamentally deal with racism in the algorithms we have, we have to center race in the models that we create.

DR. TURNER LEE: Wow. We'll come back to that, too, because I think that that's

central to the conversation around how do we begin to sort of identify where those biases show up? And then what are the mitigation strategies?

I want to, before I go to Dariely, just talk a little bit again, if you have a question, please send it to events@brookings.edu. And please, let's Tweet this conversation out, @BrookingsGov or #AIBias.

Now, so far, Dariely, and I'll need you to unmute, as well, we've heard so far around I think what Fay has sort of insinuated at, which is who's at the table at the design of these models? And that directly correlates with the number of computer scientists or engineers that are graduating that are taking that level of advanced education, right, and applying it to new AI systems.

And then I think we also heard between Fay and Rashawn context. What happens when you try to apply the context, but what's the context of the application of the model?

Dariely, I want to talk to you about the outcome, right, and the determination that actually comes from a model when you start to think about race and racism. In particular, I know that you've worked a lot on this around employment.

Now, you're a lawyer, so it's an appropriate question to you. What is your opinion in terms of this evolution of algorithmic models and AI systems that may actually come back with these types of determinations that in many respects could mirror what both Fay and Rashawn are talking about?

MS. RODRIGUEZ: Well, first, I want to say thank you, Nicol, for gathering us today on Juneteenth. I think we are living in a very important time where the issue of structural racial discrimination is on the table. And as long as I wanted to be a civil rights attorney, I've been waiting for the moment when structural racial discrimination was being talked about in the mainstream, and here we are. Right? And so I think that now is as good as a time as any to have this conversation because I think we have a real opportunity to push for change and to ensure that entities that are using these algorithms are being held accountable.

So, Rayshawn was talking about the current crisis that we're in. We're grappling with a pandemic that is ravaging our country, that is disproportionately impacting Black and Brown communities,

not only from a health perspective, but also from an employment perspective.

We are roiling in a recession and, not surprisingly, Black and Brown workers are being hit the hardest. We are seeing unemployment rates that are at historic highs. The Black unemployment rate has always historically been twice that of the White unemployment rate, but now we're seeing that the Black unemployment rate is the highest it's been in a decade. And the Latino unemployment rate is also incredibly high. We know that Black and Brown workers are the last to be hired, the first to be fired. And that's because there is a system of structural discrimination that has been long-lasting that results in overrepresentation of Black and Brown workers in low-wage jobs and underrepresentation in supervisory and executive roles.

We also know that there is a huge racial pay gap where workers of color are paid less for the same work as White workers. So, we have all of these compounding factors that are being just starkly revealed now. And now race is in the middle of the American conversation. Right?

And so, I think it is really, really important for us to talk about the role that algorithms play in creating or limiting opportunities, specifically for communities of color in this time. We know that because of the recession very few jobs are hiring. A lot of jobs are terminating employees.

And so, what I've seen in my practice is that oftentimes part of the reason that businesses turn to algorithms is for efficiency, right, for cost savings. And as you can imagine, in this environment, in this economic environment where we have a recession, businesses are going to look to be even more lean and even more efficient and to increase their cost savings.

So, I've seen some articles saying that we know that there was already a trend of businesses turning to algorithmic decision-making and hiring in the past. And that's only going to further increase, I think, in the midst of this recession. And why is that problematic, right?

I think that a lot of businesses think that if you're using these algorithms, then you're removing subjectivity from the process. Right? But oftentimes what the algorithms can do is replicate historic biases and historic discrimination. Right? If you're looking at particular industries, like the tech industry, where historically there has been an underrepresentation of workers of color, if you're developing algorithms to replicate the status quo, then you're going to end up disproportionately

impacting Black and Brown workers.

And so, I think that this conversation is really important and I appreciate that you were intentional about making it interdisciplinary because I do think that it's critical for there to be conversations across sectors. Quite frankly, right now we don't have the luxury of having any artificial barriers to economic opportunity, point blank. Our civil rights require that there be equal employment opportunity at every facet and that is more important now than in a very long time.

So, we need to work together to ensure that businesses are held accountable, that they are transparent. Right? I'm sure we'll start talking about the black boxes of algorithms and what that means. And hopefully, you know, the folks who are more technologically savvy than I am can talk a little bit about that. But it's the real issue, right, because we need to get behind those black boxes to really understand what criteria are being used to identify people in the employment context, to identify those who are qualified for jobs. And if those criteria are not job-related, they really should not be in an algorithm.

And so, I think we have a lot of work to do to just push our transparency and accountability.

DR. TURNER LEE: So, I love the way that you all have kind of laid this out. And I know that we have scientists listening, I know we have policymakers listening, and I know we have people who work in general industry or are just very interested in the topic.

Fay, I want to go back to you, right, because Dariely just laid out a series of data points. And part of the challenge is, and I know I talk about this in my work, is when you look at this continuum of bias. It can start in the question, you know, what is it that you're trying to develop around, you know, whether it's an employment algorithm, it's one that helps with financial determinations or with something on healthcare, but in the middle of that you have training data. So, what are you using to train to actually get -- what data are you using to train the algorithm to help you in reaching what your preliminary goal is? And then finally is the outcomes.

I want to stick on that training data side. That's that black box side that Dariely insinuated. What is it about the training data that I want people to know? Because I want people to not

walk away from this saying, hey, you know, I'm not explicitly being biased or intentionally being biased or discriminatory when I'm developing algorithms. But what way, Fay, does the training data sort of influence the type of determination that you actually come up with?

DR. PAYTON: Interesting question. So, you know, the role of training data, and I have to applaud my two colleagues here, and I'm picking up on this term, as Rayshawn has said, you know, centering race and design and Dariely is saying this black box of algorithms. And so, it really boils down for training data, if we think of training data I would like for people to really think about what's the source? What's the data source? What is the source of that data? How has it been sourced? What attributes does it have? What's the accuracy? How clean is that data? And first of all, can you use that data to ask or answer the right questions or the questions you're trying to get to?

Ofttimes it may be a forced fit. And so, this is why you see many people now as we try to democratize data and AI, people are starting to and scholars are trying to build their own datasets. Right? But the idea is that the training data will actually teach or helps the machine to learn to make predictions. That's the first step.

And while that's a natural first step, if that first step has some degree of variance, if we think about it, if it's a .001 in the beginning, if it's biased from the beginning, it's only going to exponentially have the same sort of bias built in on top of bias, right, toward the end. So, there's the training data. And that training data should help inform the validation data.

And if that validation model is really what is hoping to be developed to be the "right" -- and I air quote that -- what's the "right" application, to say does it really address the problem that we're trying to identify?

Now, the problem with the validation data what happens is that we tend to overfit. And when we overfit, we get instances like Rashawn has talked about, is who's being identified by the geosystems when people are protesting? Who is going to receive or not receive their approval on the mortgage lending application, right? Or Dariely said, what about employment? Who is going to get the next call based on what the algorithm -- so, that overfitting can really be problematic.

And then lastly, there's the test data, which should get to your final truth. And that final

truth piece should relate to, and I always say, it's the final truth that should relate and tie to real-life uses and experiences. And if we go back to your original question, if we don't include race in this, how do we then get to the real truth of the matter in terms of constructs, factors, and the lived experience?

And I just want to say this one thing, when we think about that training data, we should also think about is there algorithmic fairness built in? And that fairness has other implications on how you center up race in the design process. Because you're going to hit noise, you're going to get variance, and you're going to get bias. And those three have implications for Black and Brown populations.

DR. TURNER LEE: Yeah. No, I mean, I want to go to you Rashawn. And thank you for that. And everybody can just open your mikes because we're all in my living room, we're going to have a conversation. Right?

DR. PAYTON: Okay.

DR. TURNER LEE: I mean, I want to talk to you, Rashawn, about that because there's two parts of this that I think Fay is bringing up.

One is can we ever achieve inclusive datasets or datasets that allow us to see the full picture without building on top of bias? So, for example, in areas that you work on, like criminalization, will we always start as designers or developers of algorithms and interpreters of algorithms and licensers of algorithms with data that is inherently flawed before it even leaves the door?

MR. RAY: Great question. I mean, I think if we use existing crime data, yes, because crime data is flawed. It's oftentimes inaccurate and it's incomplete. And so, I think part of it is the data source, as Fay was saying.

And I think one of the key things I always think about, particularly when it comes to policing and, you know, at the Lab for Applied Social Science Research in Maryland, where I'm the director, we've developed this virtual reality program. And we have an algorithm where the virtual reality characters respond to police officers. And we had to think deliberately as we were helping the corporation we were working with to build this algorithm that we brought in a diverse group of officers, diverse not just in terms of race, but in terms of gender, in terms of region, in terms of age, in terms of experience, in terms of specialty units. So, the input was very deliberate in what we did. And we

continuously had safeguards and checks along the way.

So, as an example, when we brought some of my Southern colleagues up to our lab -- and, of course, Maryland's still Southern, but the Deep South, where I'm from -- is they would get annoyed when virtual reality characters would say something to them or respond to them in a particular way, in ways that people, say, from Philly or New York or another -- you know, St. Louis, they wouldn't react in the same way. Well, you know, that should be part of the algorithm. But if we were only relying on that department or those few departments, we would have something that was flawed.

And so, I think that the input is extremely important and oftentimes what we try to do is draw upon existing data. I think maybe it should be more part of it, but we have to identify where the gaps are.

And I think the bottom line is this. Policy is oftentimes way behind technological advancements. And one of the things that Fay said is as it relates to the questions to be asked.

And, Nicol, I mean, I hope that you share some of the wisdom you have, even though you're the moderator, because, you know, you've really helped to push legislation and helped people to think through, help policymakers to think through the legislation. So, I think the Algorithmic Accountability Act, the Justice and Forensics Algorithmic Act, I think these are put in place not even from the side of researchers, but more so for corporations who are advancing technology so quickly that policy can't keep up. The academy can't even keep up.

And then what happens, particularly in policing, because I think this is where the bias really happens. One thing I know working with all these police departments across the country, if you have a new, shiny toy, a new, shiny program, and they don't have to pay for it, they're going to normally be all-in. So, they're like, oh, okay, so we just have our officers participate and you'll let us have the technology? Yeah, I mean, that's kind of the exchange because we want the data. And they're like, okay, great.

Well, what are the problems with that? The problem is that they are now using a product that hasn't been tested, that hasn't been certified, and now all of a sudden they're going to apply those metrics and skillsets that they've learned when they go out and interact with people. And there needs to

be some policy and regulations. Even though I'm in this space and, you know, the wild, wild West might help a person like me because I think at our lab we do things pretty objectively and straightforward, but there's still needs to be regulations in place to prevent the biases that are going to pop out.

DR. TURNER LEE: Yeah, and I want to talk about that. And before I talk about that, you know, I've got to give credit to where credit is also due. There's a whole lot of people, Rashawn, in addition to me working on this. Fay is one, Safiya Noble, Rakiya (phonetic) Benjamin. There's just a -- I think what's happened is that there is more of us, and I'm a sociologist, that are looking at race from a microscope to figure out how to, like you said, improve upon the discipline.

And this brings me back to basically what you said, Rashawn, in terms of being really sensitive to youth cases. Right? Going back to what Dariely talked about in terms of the employment.

So, Fay insinuated that maybe we need to ask people their race, right? Maybe we need to put that at the center of models, and I'm actually one of those people who thinks that way. After talking to a lot of companies on facial recognition who tell me that, one, and we know this from people like Joy Buolamwini at MIT, that facial recognition technology is not good for recognizing darker skin hues. Black women who change their hair is one. And I tend to tell companies I'm going to change my hair over and over and over again, and I want somebody -- some technology that knows what I look like, so ask me. Ask me to be a part of the sample.

But you're a civil rights lawyer. And I've been putting out there for a long time, along with the Lawyers' Committee, that we should update the civil rights laws to ensure that these models are also in compliance with things like the Fair Credit Act, the Fair Housing Act, you know, that there are litigated cases, settled cases where civil rights are protected.

If we were to add race into this equation, this if my question, this is my Oprah moment, how would that impact, literally, like companies' use of proxy in place of race? And how would that look -- what would that look like if these laws were in place around compliance from a civil rights perspective, right? Do you think that we should put those guardrails in to help with better development or should we put those guardrails in to protect? Do you see what I mean? Like what value do you see modern-day civil rights laws being updated and applied to the digital space?

MS. RODRIGUEZ: Well, I guess I'll start by saying that the current civil rights laws that we have now do apply, specifically in the employment context. Right? Title 7 has always outlawed disparate impact, as well as disparate treatment, right, starting with the *Griggs* case in the '70s. And then you have the Uniform Guidelines on Employment Selection Procedures that basically say that any employment selection procedure used by employers has to be business-related. Right?

And so, the laws that we have now do apply and they do prohibit policies and procedures that result in a disparate impact against protected groups if that disparate impact is not business-related and inconsistent with business necessity. The problem is, is that the laws that we have, the Uniform Guidelines, are old. And they do need to be updated to reflect the realities of today, right, to reflect the technologies that employers are using to not only recruit, but also to hire, to interview, to evaluate. Right? A lot of employers are using technologies to evaluate productivity. They're using them to determine who should stay on the job and who should be let go. And so, you have a whole scheme of technologies that are being used in the workplace and literally affect every facet of a worker's day-to-day interaction with an employer beginning from when even before they are considered an employee.

And so, I think that employers do have a lot of information that they can use to track whether the algorithms or the technologies that they are using are having a disproportionate or harmful effect on the applicants or candidates or employees. And I think part of the issue is that oftentimes the employers are working separately from the developers. Right? And Dr. Payton got to this a little bit.

It's important for those two groups to work together. You can't just offload the responsibility of developing an algorithm or a selection procedure to a developer and have them come up with something that you need to use for the selection for your job. Right? As an employer, the employer knows what criteria, what qualifications are necessary for a job. And they have to communicate that information to the developer and the developer then has to use that information to develop an algorithm that is related to the job and consistent with business necessity. Right?

I think what happens is that oftentimes they're just trying to look for the best employee, whatever that subjectively means to them. Right? And so, if in X company that means the person who lives closest to the company or it means that, you know, there's a correlation between individuals who

play lacrosse and they tend to do very well at the business, well, guess what. That's going to replicate systemic discrimination, likely on the basis of various protected categories. Right?

And so, I think that we do have the laws in place. The laws do need to be updated. And businesses need to do a better job of actually auditing the systems that they're using to determine whether they are having a disparate racial impasse. And if so, to ensure that they're going back to their developers to say, hey, wait a minute, this needs to be fixed.

My concern is that before, right, when we had the Uniform -- when the Uniform Guidelines were created, they were created to get at things like tests. Right? Like physical ability tests or criteria like educational testing, right, that employers sometimes used as a pretext for discrimination on the basis of race, on the basis of gender, to perpetuate segregation.

Now, those have been replaced with these more automated systems. Right? And when you're using these technologies, you are really amplifying the capacity of your HR team. Right? You're increasing the ability with which they can sort through hundreds and thousands of applications. And so, there's a real opportunity for turboing discrimination, for amplifying the discrimination. And that could potentially have very grave implications that disproportionately impact many more people.

DR. TURNER LEE: You know, 20 years ago, and Rashawn probably remembers this, there was work in sociology around people who didn't get employment because of the sounding of their names or the location or the address of their resume. And this was almost 20 years ago, we're talking about how people get work. And it's so interesting to me when I'm in this area about how many of those trends from the sociological field sort of play out in terms of the outcomes that we actually see coming from algorithms.

Fay and then I'll go to Rashawn on the same question. So, what do we do about it, right?

DR. PAYTON: Yeah.

DR. TURNER LEE: I mean, if we know that there's a certain sensitivity when it comes to heightened vulnerabilities around race and I think what Dariely's explaining so well around racial discrimination based on disparate impact, although, Dariely, we can talk about this when we have more time. It's so hard to prove that in a very opaque online system because you're not necessarily applying

for an online application next to somebody else. Right? We would say, oh, like the Apple card. I'm applying with my husband. He got it, I didn't get it. How many people of color know that they're not seeing something or they're not, you know, eligible for something?

So, Fay, what do you do about this? Like how do we actually move the needle so that we can, I think, bring what you all are trying to say at the table? And we'll take questions in just a few minutes, folks, so please email your questions or post them.

What are the best practices then to be very conscious that race is factoring into the design models and, you know, subsequently, can show up as the algorithm itself is framed within a certain context?

DR. PAYTON: Yeah. Can I give you my laundry list? (Laughter)

DR. TURNER LEE: I mean, people want to know this.

DR. PAYTON: Yeah.

DR. TURNER LEE: I mean, all the people outside that want to know this.

DR. PAYTON: So, I think the question is what are some promising practices? Because there always needs to be some degree of human intervention, let me say that. Right? So, I think one of the first things I would say is there is always this focus on some professor of information technology and analytics, right? So, there's always this need to have big data. Big data, big data, big data.

We need to listen to small data. Small data tells you a lot about the lived experience. Small data tells you that I'm not healthcare-adhering to my diabetes medication not because I don't want to, but because I don't have transportation to get to the medication. I mean, those are really, really fundamental questions.

In sort of this tech data space, that means we need to look at and analyze the meta data. So, small data matters. That small N, that small N matters.

I think we published this paper on looking at -- some colleagues and I published this paper on looking at mental health among college students. And we used natural language processing and we built this great corpse of papers that looked at it and we came up with this model. And the bottom line is we needed to hear what the students themselves had to say about mental health. So, small data

matters.

I think this idea of we talk about small data, we talk about these proxies. You know, there's some part of me that says that how can we reduce the impact of the proxy effect? And the idea of the proxy effect, as my colleagues have so notably mentioned here, we got ZIP Code. What does ZIP Code determine? Well, if you use me as an N of 1, ZIP Code would have told you I wouldn't have been on this panel with you guys at the Brookings Institution and would not have been a Ph.D. Right?

So, I think we need to look at this as an ecosystem, not just what research question are we trying to answer, but is what we're trying to answer, how does it play into the broader ecosystem by which people have to navigate? There's housing, there's employment, there's healthcare, there's policing, there's talent management.

I do think that, yes, algorithmic audits are critical. I think the algorithmic audit is what happens. Can we do an audit? That's where the researchers found that, hey, yes, there are high costs, but Black and Brown folks aren't necessarily receiving the same quality of care and the care was withheld. So, the algorithmic audit should always ask the question as a promising practice, how did the audit come out, but who did the audit fail? Not looking for success and looking at how did it fail? And if it failed, then why did it fail and who did it fail?

So, I think I'll stop there and give others an opportunity to answer. But I think those are some things that we need to talk about, in addition to ethics.

DR. TURNER LEE: Yes, yes. And we're going to actually end on ethics before we actually go into the conversation.

Rashawn, what do you think? I mean, how do we address this from a policy perspective? I'm getting a lot of questions, I'm not sure how many I'm going to be able to answer, but I want to integrate this into your response, as well. From an academic perspective, are we doing something wrong? You're in the academy part time, too, right, as Fay? What should we be teaching people, as well?

MR. RAY: I mean, I think, first, we need to just keep listening to Fay. (Laughter)

What I think in addition to that, you know, I think it becomes about, again, the questions

we ask. So, a couple of the things that I've tried to do is to lay out the questions policymakers should ask. One of those questions includes what safeguards should we put in place to ensure that technology is being used properly and it's doing what it's intended to do?

So, when it comes to policing that might look like a moratorium on when the technology is applied and how long you can actually use it for. And I think in law enforcement it's a good example because we already have a lot of regulations and laws in place that guard against things like what happened to Breonna Taylor, for example, which I'm very -- like one thing I'm always interested in when it comes to no-knock warrants is how do you so easily always go to the wrong house? Like I just couldn't -- even studying policing, that is something that I always come back to.

And thinking about that, some of the things that have been put in place, so, of example, the Facial Recognition Technology Warrant Act was put in place to essentially put a 72-hour safeguard in place. You've got to get a judge's approval. I think for law enforcement that makes sense.

I mean, if we're talking about companies, well, what does that look like? Well, I think similar to academic research standards, companies should have to produce results at the .05 level. Other people say, no, it needs to be extended, particularly when we talk about big data. Then it needs to be at the .01 level or .001 level. So, for people trying to think about what I'm saying, I mean, we're talking about statistical significance where we can essentially say within a particular confidence interval that 95 to 99 percent of the time that this is happening, meaning that there is no bias that's existing there.

And so, what needs to happen is that companies could partner up, say, with law enforcement agencies or companies could partner up with research agencies, and actually carry out some of these processes.

And I want to highlight something Fay said when she talked about big data and small data because that's one of the things that we do in my lab. As we built this algorithm, we had these officers go through these virtual reality scenarios, I mean, which is super neat. If anybody has kids and maybe playing "Fortnite," that's kind of what it looks like. It's super cool. Police officers love it. But were changing who they're interacting with in a setting. So, we varied the race of the person. They go from Black to White or they go from White to Latino. But the person talking sounds exactly the same. It's the

same person's voice. We just tested their skin tone and changed their skin tone. We do the same thing as it relates to gender.

And so what's interesting there is we had them go through all these scenarios: traffic stops, domestic house scenes, robberies of stores. And we noticed that this one group of officers, I mean, they were just really getting upset. And they started pulling guns on the virtual reality characters who -- and the characters get loud. They get loud, they start cussing at police. I mean, they escalate. The algorithm is built to respond based on what the person says, what the officer says, how close they are, and their tone of voice.

So, we started noticing that these officers were shooting or pulling their guns on the virtual reality characters. We were like, whoa, what is going on here?

The small data was we do a debrief interview afterwards. And what we found out was that this group of officers had worked nights, the night before, they hopped no a plane at 6 a.m., flew to D.C., hadn't had a lot to eat, hadn't had a lot of sleep, and then they get in this virtual reality scenario and they are short-tempered, upset, hungry, and stressed. You give a person a gun, that's a lot of times what happens when it comes to officer-involved shootings.

But if we didn't have that small data, and I think this is what's key about building algorithms, we take those small data and then we create variables out of them to put into our statistical model. So, we can take did they sleep or did they not sleep? Did they work the night before? I mean, a series of additional variables that we didn't know at the time that we were supposed to collect, so then we can put it into the model to see what's explaining that.

It explained away the regional effect. It explained away the race effect. Like we don't find racial differences in the likelihood of officers using use of force on people. That's extremely important for people to know. I think algorithms can be used to help with that. But if you have a flawed model, it will give you flawed results.

DR. TURNER LEE: You know, I like what you're talking about and I'm going to end with Dariely and then we're going to go to questions. Because what you're basically suggesting, and this is something that I'm doing in this model that I'm building around an Energy Star rating of algorithmic

fairness and efficiency, is basically combining self-regulatory public policy and civic engagement all in one, so we can actually, you know, have people respond to whether or not the algorithm is correct; have companies sort of figure out -- tell us how they're testing these models. And then in the middle, you know, maybe there are, like Dariely was saying, some pressure points for public policy to come in to make sure they're lawful, et cetera.

But what you're saying, and one part of my model on the self-regulatory side, I make the argument that we need companies to try a whole lot of stuff. Right? And what you're basically saying is, and Fay is saying, small data matters as much as big data, in addition to really testing those algorithms on primary, secondary, tertiary conditions to determine which ones can actually generate less bias. And I think that is so important.

Dariely, I'm going to come to you because I think what we're dealing with in technology, folks, is permissionless forgiveness. We used to have permissionless innovation because we were so excited about all these new developments, but now we're seeing companies who are trying to be first. And when they're first and they make a mistake because they haven't tested in different contexts, they tell us that we're sorry. But the problem is that's one more person who didn't get approved for a loan. That's one more person who was detained in the criminal justice system. That's one more person who may not actually get that employment.

So, Dariely, what I would ask you is, do we need government to come in? Is this a government issue or do you think what you're hearing from the sociologists and the scientists are maybe some ways to sort of mitigate some of those risks associated with race?

MS. RODRIGUEZ: Well, I want to build off what you were saying before after my last response about the fact that there is a lack of transparency. And oftentimes if you have been discriminated against because of some sort of algorithmic decision-making model, you don't necessarily -- you don't know. Right? How can you know that you were the subject of a negative employment decision because of some sort of technological model? There's just no way to know, especially if it's happening on the front end, at the recruiting or the hiring process. Right?

And so, I think that because of that lack of transparency that exists -- well, A, I think

employers need to be more transparent and let their candidates know exactly the processes that they are going through. But, B, short of having that level of transparency, governmental agencies just have to step in. Right? Because employment agencies in particular, and I'm thinking about the Department of Labor, the Office of Federal Contracts Compliance Programs, it's a mouthful, but they are the agency at the federal level that is responsible for ensuring Affirmative Action obligations for federal contractors.

We know that businesses who receive federal monies, they employ a large swath of the American workforce. And when they receive federal monies, they have the obligation to ensure that their polices and practices are non-discriminatory. And they have to regularly audit them.

So the Department of Labor has the ability to and does go in and audit these businesses, right? And the positive to that model is that they are able to go behind the scenes and really understand the practices and policies that employers are using to hire, to promote, to terminate. And they're able to get access to information that individual job-seekers don't have. Right?

And so, I think that the government enforcement agencies, like the EEOC, like the Department of Labor, do have a special obligation here to prioritize this issue and to ensure that they are dedicating the resources to investigating potential discrimination happening in this space. I think the issue we need to ensure that resources are being allocated to those agencies through the appropriations processes.

And then we also have to ensure that we as advocates are holding those enforcement agencies accountable to ensure that this is a priority for them.

DR. TURNER LEE: That's right. That's right. And I think that that's actually a good way to look at it because what it helps you to do is to figure out what parts of the government, right, can you empower to sort of help solve this problem?

I want to go to some questions and I have a lot, so I'm going to try to combine them. So, pardon me, listeners and viewers, if I actually mess this up.

But I think one question that I'd like to ask of all of you and anybody can respond, and then we'll go to the next question, someone asked how do we actually create more of these types of conversations across disciplines and push these conversations to the product development side of tech

companies, as well as to universities? So, how do you actually simulate or bring these types of interdisciplinary conversations on race and racism to product development teams at private sector companies and to universities?

Anybody can answer. And I'm going to try to sort of firehose these because there's a lot of questions, so if I can get one person to answer, then I'll move on to the next question.

MR. RAY: I mean, well, I'll just I think at universities we have to de-Balkanize our reward structure. So, you know, in every discipline, no matter what it is, the point is to publish, but it's typically to just publish within your discipline. It leads to us not looking outside of our discipline for references, for citations, for articles, for collaborators. And I think that's something that has to change for the academy, but also at universities, that they can start to rethink this.

So, I know at University of Maryland, the one thing they do since I do a lot of interdisciplinary work, you can actually come up with a separate document that you come to terms with what's your department and your college for which journals and which types of collaborations are going to matter. So, in sociology, Nicol, as you know, I mean, kind of being the only one, or the *numero uno*, is a big deal, whereas in other disciplines, whether they be in health or research and science fields, the point is to have a series of authors and maybe be the last one or the first one, depending on where you are in your stage.

So, I think that those type of regulations and kind of the decoupling of that would help people to collaborate, and then providing resources. So, Maryland has done a good job at this. They actually have these grants where you can only get the grants if you're collaborating with people in another discipline and preferably outside of your college. That kind of forces people to collaborate.

And that's kind of what led to our virtual reality program. We collaborated with some computer scientists. We put all of our graduate students in a room together over our winter quarter. We bought them a whole bunch of food and they created something amazing. But it came from the university kind of having opportunities for people to collaborate across disciplinary divides.

DR. TURNER LEE: Right. Now, Fay, I have a question for you. You brought up the ethics piece. I'd love for you to sort of unpack that.

Do you think a conversation around fairness and ethics would actually get at what we're talking about with regards to race and racism?

DR. PAYTON: I think it would.

DR. TURNER LEE: (inaudible) question.

DR. PAYTON: Yeah. So, I think that it will get at it. I think -- but there's much more to do along the fairness, the equity, the accountability, the transparency piece. I think the conversations need to be inclusive, particularly you mentioned something else earlier in the conversation, that it's not just you doing this work. We are talking about an interdisciplinary team here, but I'll say there are Black and Brown women that are doing the work.

Now, whether corporate, industry, academe puts credibility on that or listens to what is being generated by these women, I think that piece is really important and should not be ignored because they really are doing great work in this space. So I think the ethics, yes, is important, but I'd like to push the conversation a little further, so that we can get to our solutions.

And the last thing I will say about that is, you know, we have in this sort of academic space this idea of there is bench science and then there's translational.

DR. TURNER LEE: Right, right, right.

DR. PAYTON: Both are important, but not to everyone. And so, where ethics is involved in each side of those siloes without being siloed, as Rashawn has just talked about, is very critical to that piece.

And look, industry doesn't get a pass here. There is a lot for industry to do in this space and that does include the women that we are talking about and all those other scholars, as well.

DR. TURNER LEE: Yeah. I mean, I try to -- we did a paper at Brookings last year, we used this concept of diversity in design.

DR. PAYTON: Yeah.

DR. TURNER LEE: And what you just mentioned, I would love to add to the national debate on AI and ethics: inclusiveness, fairness, ethics, and equity.

DR. PAYTON: Yes, definitely.

DR. TURNER LEE: That's going to be -- for those of you who know me, I think that really rung a bell with me in terms of how we're thinking about these results.

Dariely, I'm going to give you this question. As a lawyer to somebody who follows privacy, from Lauren, under most privacy laws race is a sensitive category or a special category of data, meaning there are restrictions on collecting it. I think this might be a barrier to inclusive and equitable design. Do you agree? If so, can you share your thoughts about how we would address that?

Because that's an interesting question, right? How od you actually develop these systems if privacy is still a number one, fundamental right of individuals?

MS. RODRIGUEZ: So, it's interesting because in the employment space, most employers that are subject to civil rights laws do need to collect race, ethnicity, gender information from their applicants. And then they also have to solicit that information from applicants. And they have to collect that data from their employees.

So, I think for employers it's pretty unique because typically they do have quite a bit of data to, at a minimum, audit these systems to see what results they are having. So, that data is being collected and employers should be using that at least on the backend to ensure that the systems are not having a disproportionate impact.

DR. TURNER LEE: Yeah. And I want to say, also, to that question, some of you may have seen there was some recent announcement by Airbnb I just looked at the other day and I think I shared it with the group around coming up with independent research studies to actually help with some of that, particularly if you cannot put race in because of privacy constraints.

I want to say because we're running out of time, but I do have one last question I think is really worthwhile to ask on this day of Juneteenth. I want to say for those of you who are listening and who are also asking about questions around regulatory or policy solutions, please go to the work that we're doing at Brookings. We have an AI initiative that takes into consideration governance, bias, and national security.

If you go to the Brookings website you'll actually see a series of papers that look at regulatory and policy concerns, in addition to issues around fairness and discussions around

transparency. Some of the research is my own and some of it comes from some notable scholars that are working on this stuff around the country.

I want to end with this because it is Juneteenth. And as we started this panel, when the people in Texas were actually -- you know, found out that they were actually liberated and free, they celebrated. And that's why we celebrate this today.

Khalid (phonetic) gave a really good question. Because oftentimes when we look at race, we look at the glass half empty. His question is, I have consistently only heard about the failures of AI models to further equality and equity. But if there any research that actually tells us that AI frameworks and models can be used to intentionally reduce inequity or uplift groups of color?

And I want to end there. Let's end on a note with which we can actually look at a solution. Anybody know of anything or anybody imagine ways in which AI could be used to actually change the equation of how it's currently being, you know, sort of diffused?

Boy, that's a long silence. (Laughter)

MR. RAY: I mean, one thing I would --

DR. TURNER LEE: I know, look, I know we've been mad, Rashawn and Dariely and Fay, but let's see if we can find a glimpse of hope.

DR. PAYTON: Glimpse of hope. Glimpse of hope, okay. Okay, okay.

MR. RAY: Yeah, I mean, I'll just quickly say, so, I think there -- so, one big thing that algorithms are being used for now is to look at bail hearings and also look at who gets paroled. And I would say that those particular algorithms and the way that researchers are combing through that holds some promise because they are really comparing the models. They are essentially doing a meta analysis of the models, and some of them show a lot of promise, particularly because they take into account race. Because if they only take into account place, then all of a sudden you're having a big bias and it's going to be bias mostly toward Black people, who might live in neighborhoods that are overpoliced and have higher levels of crime.

So, I think that that is a promising place. There have been some algorithms that have been put together that have been able to predict certain outcomes well and even predict things a little bit

better than judges' decisions.

DR. PAYTON: I have --

DR. TURNER LEE: Fay?

DR. PAYTON: Yeah, I have one.

DR. TURNER LEE: You're going to give me one?

DR. PAYTON: I'll give you one. So, when I think about AI-powered alerting systems in healthcare, right, particularly for elderly populations or populations that may need medication on a timed-dispensed sort of time period or continuity, I think those are really places and spaces where AI has done some amazing things. And also, when I think about in the terms of preventative care, there may be some -- there are some cases there where AI has been quite transformative.

DR. TURNER LEE: And, Dariely, your hope or promise as we get ready to log off?

MS. RODRIGUEZ: I think in many jobs, including professional jobs and up, typically we have historically or employers have historically relied on criteria such as having graduated from a particular school, looking at Ivy League graduates, and other criteria that don't always translate to success in a job and can be rife with biases. And so, I think technology probably has the promise of not over-relying on criteria like that, but relying on criteria that actually shows that a candidate is able to do the job.

But I think in order to achieve that promise, employers just have to be incredibly intentional and they just have to really, you know, slow down and be deliberate about ensuring that the teams that they're building to develop these models are diverse; that within their own businesses that their executives are diverse; and that the criteria that are going into the black boxes, if you will, are rarely related to the jobs.

DR. TURNER LEE: So, I will say thank you for that. And, I mean, I would add, also, just one snippet at we wrap up.

First of all, AI can be harnessed for good. Right? But I think all three panelists, and I thank you all for being here today, you've got to put intentionality up front and you have to bring some level of interdisciplinary cooperation and, at the end of the day, you have to be aware.

And so I thank everybody for actually joining us today because, guess what, right now we are very aware of what's happening around us in terms of these very difficult times. And the issue of race and racism have probably hit every dining room table and every coffee shop -- well, virtual coffee shop because we're literally working from home, but it's hitting the public domain. And I would say to all of you keep letting it hit it. Because if we're going to create AI that is equitable, fair, inclusive, and doesn't have that tradeoff, and it's going to take conversations like this.

So I want to thank each and every person that chimed in. Again, thank you, panelists, for being here. And at Brookings we're working on this stuff. Please stay in touch, keep the conversation going on Twitter, and happy Juneteenth.

MS. RODRIGUEZ: Thanks, Nicol.

DR. PAYTON: Thank you.

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

Carleton J. Anderson, III

(Signature and Seal on File)

Notary Public in and for the Commonwealth of Virginia

Commission No. 351998

Expires: November 30, 2020