THE BROOKINGS INSTITUTION

WEBINAR

SHOULD THE GOVERNMENT PLAY A ROLE IN REDUCING ALGORITHMIC BIAS

Washington, D.C.

Friday, March 12, 2021

PARTICIPANTS:

Moderator:

ALEX ENGLER
David M. Rubenstein Fellow, Governance Studies
The Brookings Institution

Featured Speakers:

GHAZI AHAMAT
Senior Policy Advisor
Centre for Data Ethics and Innovation

NICOL TURNER LEE
Senior Fellow and Director, Center for Technology Innovation
The Brookings Institution

LARA MACDONALD
Senior Policy Advisor
Centre for Data Ethics and Innovation

ADRIAN WELLER
Programme Director for AI
Alan Turing Institute

* * * * *

P R O C E E D I N G S

MR. ENGLER: Good morning. Good morning, hello everyone, and welcome to this Brookings Institution panel on the role of governance of algorithmic bias.

My name is Alex Engler, I'm a David M. Rubenstein fellow in Governance Studies here at the Brookings Institution. I focus largely on artificial intelligence harms. And one of the best known, what's important to harms is algorithmic bias, something we hear about all the time.

And that's what we're going to talk about today, specifically what is and in what situations does government come into play. In the models themselves there's really no doubt that bias can manifest. We've seen models of human language that have demonstrated bias against women and people with disabilities, speech recognition models that have demonstrated bias against African Americans as well as problems with different dialects and regional variations. And certainly issues with facial recognition has performed worse on people with darker skin as well as intersectional groups like Black women. Unfortunately I could go on for quite some time the amount of research about these problems is really quite extensive.

Sometimes these biases are found in just the models, but they've also been found in the deployed commercial products and the websites and in the government services out in the world. Again, unfortunately, countless examples, provisioning healthcare and providing job opportunities and directing the attention of the police in determining credit worthiness. We've seen well-documented convincing evidence that these algorithms are at least continuing the biases that pre-dated them. Still the reality is that it's hard to tell if a net were better off or worse off.

Certainly the algorithms have bias. It can be difficult to know if it's definitely worse than it was before. Many of these biases are not new, the algorithms are replicating them. Certainly there's evidence that humans are also biased in their decisions around employment, policing, finance, etc.

At the same time you could make a compelling argument that these algorithms are taking that bias to new scales, right, when an algorithm works effectively it's not really doing a small number of people, its performing its service for thousands, tens of thousands, even millions of people.

However, in some cases active interventions by AI developers and also building systems, sociotechnical systems around those algorithms can make them better. There's an excellent conference just this week on AI fairness, called The Fact Conference. And then that field of responsible, ethical algorithms field, is making a lot of meaningful progress.

Yet again, despite the fact that there is the possibility to make the algorithms more fair. We don't really know if that's happening out in the real world. It's expensive and time consuming for expert data scientists to spend their time building more fair models when they could be building a new feature to sell or when they're focused in a different area to political consideration.

All of this combined is a pretty complicated overlap of possibility but also potential harm. And we're not really sure by default, we shouldn't feel confident by default that the world's going to be better off if we just let algorithms start influencing large and critical areas of the economy and progressively more in the public sector as well.

So that's what we're going to talk about today. I'll be moderating the conversation. Excited to have my colleague here, Nicol Turner Lee. She is a senior fellow in Governance Studies and the director of the Center for Technology Innovation of Brookings Institution.

We are also joined by three guests from the Centre for Data Ethics and Innovation, which is an independent review board commissioned by the government of the United Kingdom to discuss its emerging role in algorithms.

Specifically I'm joined by Ghazi Ahamat, who is a senior policy advisor at CDI, as well as Lara MacDonald, also a senior policy advisor there, and Adrian Weller, who sits on the CDI Board as the programme director for AI at the Alan Turing Institute and also a principle research fellow at the University of Cambridge.

So in just a second I'm going to hand off the baton to our guests from the Centre for Data Ethics Innovation for a presentation of their work on how this is playing out in the United Kingdom. Right before I do that I just want to remind our viewers you can submit questions to the speakers by emailing events@Brookings.edu or via Twitter using #AIBias. So if you want to submit a question during the talk, again that's events@Brookings.edu or via Twitter at #AIBias.

Again, thank you to our panelists at CDI for joining us. And please take it away.

MS. MacDONALD: Thanks very much, Alex. I'll just share my screen, that should be working. Well thanks very much, Alex and to Brookings for inviting us to speak today about our reports. And as Alex said, the Centre for Data Ethics and Innovation is an independent advisory body set up by the U.K. government in 2018 to investigate and advise on the ethical challenges with data driven technology. We work at cross sectors, often assisting partners to develop responsible approaches to the use of data, and we also do work on crosscutting themes in data ethics.

Today we'll be presenting our review into bias in algorithmic decision making, which we formally began in March, 2019 and published in November 2020. We could take up and task the CDI to draw on expertise and perspectives from stakeholders across society to provide recommendations on how they should address the issue of algorithmic bias.

We focused on areas where algorithms have the potential to make or inform a decision that directly affects an individual human being. We've typically focused on areas where individual decisions could have a considerable impact on a person's life.

We took a sectorial approach because we knew it would be important to consider bias issues in specific contacts. We chose two public sector uses of algorithms in policing and local government and two predominately private sectors uses in financial services and recruitment.

These sectors were selected because they involve significant decisions about individuals and because there's evidence of both the growing uptake of algorithms and historic bias in decision making.

And by focusing on private and public sectors examples we've been able to develop cross sets of recommendations for tackling this issue at a wider societal level.

In the review we make recommendations to the U.K. government and set out advice for U.K. regulators and industry which aims to pull responsible innovation and help build a strong trustworthy system of governance.

It's also worth noting that once we did look at international examples of practice our focus was really on U.K. examples and what could be addressed through U.K. legislation, regulation, and governance.

Ghazi and I will now run through some of the key findings and recommendations of the review which we hope will help frame a decision and response to the overall question of today's event.

Adrian will then make some closing remarks and we'll hand back over to Alex for the panel discussion.

So the question we're looking at today is should the government play a role in addressing algorithmic bias. Our review looks at two key ways the government can address algorithmic bias. These are the government and the public sector as a major user of technology should set standards, provide guidance, and highlight good practice.

And secondly, that government as a regulator should adapt existing regulatory frameworks to incentivize ethical innovation. But before going into these two areas, I'll set out how I reviewed to find the issue of algorithm bias.

We know that the growth and the use of algorithms has been accompanied by significant concerns about bias. In particular that the use of algorithms can cause a systematic skew in decision making that results in unfair outcomes. Our review looked at how bias can enter the algorithmic decision making system at various stages, including through the use of unrepresentative or incomplete training data, through the reliance on flawed information that reflects historical inequalities, and in failing to recognize or address statistical biases.

We also looked at the concepts of discrimination and fairness. We highlighted that some forms of bias will constitute discrimination under U.K. equalities law. Namely when bias leads to unfair treatment based on certain protected characteristics. There are also other kinds of algorithm bias that are non-discriminatory but still lead to unfair outcomes.

On fairness, this is much more than the absence of bias. Instead, fair decisions need to also be non-arbitrary, consider equality implications, and respect the circumstances and personal agency of individuals.

There are also multiple concepts to fairness to be considered, some of which are incompatible and can be ambiguous. In human decisions we can often accept ambiguity and allow human judgment to consider complex reasons for a decision. But in contrast, algorithms are unambiguous. Ultimately it's critical that the overall decision making process is fair, not merely that algorithms are unbiased.

So how should organizations, including government and the public sector, address algorithmic bias? We tried to be as practical as possible in our review and that we reach a set of suggestions for leaders and boards or organizations to strive forward. These included understanding the capabilities and limits of algorithmic tools, considering carefully about the individuals who will be fairly treated, making a conscious decision on appropriate levels of human involvement, putting structures in place to gather data and monitor outcomes with fairness, and understanding your legal obligations in carrying out the appropriate impact assessments.

Achieving this in practice would involve doing things like building internal capacity by forming multi-disciplinary teams and improving workforce diversity, understanding risks advice by gathering and analyzing data to measure potential bias, and ensuring there is clear organizational accountability and transparency.

One of the sectors we looked at in this review where the government and public sectors were using algorithms is in policing. The sorts of algorithms we found in use in the U.K. range from tools for predictive mapping, data scoring, and individual risk assessment.

However overall we didn't find these were used in a wide-spread way in the U.K., and some were under consideration rather than active deployment.

We also set out some practical advice to police forces which included conducting an intenerated impact assessment before deploying algorithms, and ensuring they had appropriate rights of access to algorithmic software. We also highlighted that national regulators should be able to order the underlying statistical models if needed. For instance, to assess risks of bias and error rates.

We also made an overall recommendation to the U.K.'s home office, which is the U.K. government's leading Department of Policing Policy to provide clarity about who was responsible for ethical oversight in governance and police use of algorithms in the U.K.

Another important way that the government can address algorithmic bias is through driving greater transparency in the use of algorithms. We made a specific recommendation to the U.K. government to place a mandatory transparency obligation on all public sector organizations using algorithms that have a significant influence on significant decisions affecting individuals. We think this would help drive the clearer articulation of benefits and risks in the use of algorithms, uphold democratic accountability, and have an impact downstream on the supply chain in private sector. Ultimately we see this as a key ingredient building trustworthiness in government's use of algorithms.

We are also very pleased to be working now with the U.K. Government Digital Service to help them develop a new approach to algorithmic transparency in the public sector as part of the U.K. National Data Strategy.

So returning back to the question, I'm going to hand over to Ghazi now who will take us through how government as a regulator should adapt existing frameworks to incentivize ethical innovation.

MR. AHAMAT: Thank you very much, Lara. I just want to briefly unpack that second point of as a regulator government needs to adapt existing regulatory frameworks to incentivize ethical innovation.

So the points that I want to talk to today is that there are existing regulatory frameworks that need to be adapted as opposed to creating this as a blank problem.

And secondly, the challenges of achieving ethical innovation. Both of those are important and the real challenge is striking the right balance between ensuring the responsible use of these technologies and unlocking the potential benefits that these technologies also bring.

If you turn to the next slide, please, Lara. As I mentioned, in the issues of bias, and particularly in the issues of discrimination, we do have existing regulatory frameworks that apply to AI.

In the U.K. we have equalities laws, which is particularly the Equality Act, and also human rights law, which sets out principles of fair treatment that apply across many different sectors. The equivalent to that in the U.S. is discrimination law and civil rights law. We can unpack some of the differences and similarities a little later in the discussion.

In addition, there are some provisions under the General Data Protection Regulation, GDPR, and the Data Protection Act, which set out fair treatment from individuals from a data processing perspective. And all of those do already create obligations and regulations about fair treatment in the use of these technologies.

We also see that there are many sector specific regulations which set out what fair conduct and fair treatment look like, ranging from in financial services to consumer rights to policing. And we think that all of these do apply to the use of AI by these sectors.

The challenge is how to update or the enforcement and the application of these regulations to the rise to algorithmic decision making.

In our review we did conclude that at this point, at this stage, there is no evidence to suggest that we need to be creating an entirely new regulatory regime for this particular problem. We think that this significant scope to update these regulations to reflect the challenges that algorithmic decision making has raised.

I'll now turn to one sector that is of particular interest where these issues have been raising, which is in recruitment. Now if we think about recruitment of the process all the way from advertising the jobs through to selecting, there's an increasing use of algorithmic decision making across that entire process, ranging from deciding how to describe a job in order to attract the right applicants, targeting ads to the right places, through to screening, whether it's screening of CVs or psychometric tests and the like. And even an increasing use of AI in the interview process where video and face recognition and voice transcription is increasingly being used during the interview process. And finally, there are steps in the final steps of selection around background checks.

We're seeing increased use of algorithmic technologies across these areas though the main concerns that have been raised around algorithms during the two middle steps, the screening and the interview step, where there's been a lot of increased use of these technologies and also increasing concern.

If we turn to the next slide. I then want to turn to one of the issues that has been of increasing complexity in applying the regulatory framework to the rise of algorithmic decision making. And one, and this is about data around protected characteristics. Many of these characteristics that people are worried about in discrimination, things like race, religion, sexual orientation, are also sensitive data that people may have concerns about from a privacy perspective. We think that however the concerns around privacy are probably bigger than is warranted by at least one more, because at least under U.K. law, collecting this data for the purposes of monitoring for bias and discrimination is explicitly allowed in data protection law.

That's not to say that the general public and customers would be comfortable with repeated collection and use of this data, so there are some legitimate concerns that need to be worked through. We also think that perhaps there is a concern that this is raising issues of bias that is easier left ignored.

We overall do think that greater use of the collection of this data to monitor for bias and discrimination is very important though we are actively looking at how we can be more innovative rather than repeatedly collecting and storing this data.

If we turn to the next slide. As Lara mentioned as an important issue of meaningful transparency in both the public sector that is a different sort of problem in the private sector. There are many different players, whether its regulators, end users, procuring executives, who need to get some sort of meaningful transparency over how algorithmic systems work in order to make a judgment of whether they should trust it.

And we think that there is a range of mechanisms that have worked in other environments, ranging from ordinance certification, accreditation, impact assessments, that could be applied to improve the transparency of the use of algorithmic systems. However we think this eco system is very latent and we've actively looking at what it's going to take to build the maturity of these services.

That gives an overview of some of the regulatory issues and levers the government has to potentially tackle the issues of algorithmic bias in discrimination in the broader economy.

I'd now like to turn to Adrian to reflect on what this might mean from the perspective of those building and applying these systems.

MR. WELLER: Thanks, very much, Ghazi, and I'd like to thank Ghazi and Lara and all the CDI team for all the excellent work they did on the bias review. I'd also like to thank Brookings for organizing this and I'm really looking forward to Nicol's comments soon.

So as we think about what ethical innovation really means, as many of you have probably seen, we've got over 100 sets of high level ethic principles that various institutes have published. And the good news is that that's a lot of overlap with a lot of emphasis on topics like fairness and transparency. And yet there's still, I suggest, really a lot of lack of clarity over exactly what these terms mean and how to deal with the inevitable interdependencies and tradeoffs which can occur between different desirable properties and performance in practice.

So I'd like to suggest the framework that I hope can enable us to move forward constructively from those high level epic principles to practice. And I suggest that to do that we need to focus on concrete settings in a way to understand exactly what things like fairness or transparency really mean and examine the way they can trade off against each other. As examples you can think about policing or hiring practices.

And the three things which I'm identifying here are one, we need to identify requirements. So to do that we need to listen carefully to users and stakeholders' various hopes and fears in order to figure out what's needed for a specific given context.

The second thing we need to do is, bottom left here, we need to build in those requirements, build them into our algorithmic systems.

And then on the bottom right what we're showing there is that we need to be able to check and enforce the requirements, as Ghazi was saying.

And of course this is an over simplification, there's no point checking for something which can't be built. And when we think about what requirements are needed, we need to be informed by what's possible to build. And indeed different subgroups may have different requirements or views as to what's fair, as Lara noted.

So we need to keep iterating between these three tasks over the life cycle of symptom design, development, and deployment.

My work personally focuses on the blue box at the bottom left. But it's really critical that we all contribute and work well together to ensure that all our diverse voices are heard as we try to build other algorithm systems which really will work for all of society.

Thanks very much. Now back to Alex.

MR. ENGLER: Great. Thank you all so much for that presentation. I'm really excited to get more into it. If you are watching and are interested to read about this more in depth, the easiest way to find the CDEI report is to google "CDEI Review algorithmic bias" and that'll probably get you there, depending on how localized your Google searches are. I think that will work though so I encourage you to check out that resource, it's a really excellent paper.

Adrian, I'm going to stick with you for just a second. Can you lay out for our audience who might be a little less familiar with the differences or similarities between the United Kingdom and the United States, maybe some core overlap or any distinctions we need to be making between the two as we reason about government oversight of algorithm systems?

MR. WELLER: We would be speaking, there's actually a lot of similarity. So I think this is a great topic for us to discuss. If it's okay I might suggest that if we hear from Nicol and her comments first, that might enable us to more easily compare and contrast a little bit, is that okay?

MR. HEALY: Yeah, of course, that is in fact the next question I was about to ask.

Nicol, especially focusing maybe first on the area of discrimination in policing, how do you think the U.S. should view this problem? Does it seem similar? For instance the CDI friends mentioned that algorithms were not widely used in the U.K. Is that the same here, or do you see this differently?

MS. LEE: Well, thanks, Alex, and good morning, good afternoon to everybody who is watching. And again, once again Brookings is very happy to have this conversation. It's just a shameless plug for those of you that want to learn more about a Brookings work, you can go to our AI Bias page on the Brookings website. We have a series of papers on governance as well as bias that are written by people all over the world.

I'm really excited about this conversation because I've been carefully following this model, so this is sort of a lifelong dream, and Alex Engler, my colleague, has been working really carefully I think on many of the use cases that we're going to discuss today.

So, Alex, before I answer your question I do have to lay out I think what Adrian has talked about, which is the fundamental premises of where there are similarities and differences in the U.K. approach to algorithmic bias and the U.S. approach.

First and foremost, I really do appreciate the fact that the Institute has really taken a careful look at what do we do when we set standards around algorithms generally. And when we think about the elusiveness of fairness and how we determine what tradeoffs we're actually making within the ecosystem.

For example, for every time that we think a predicted decision is actually close in terms of its preciseness, what tradeoff is being made to a human subject that may not be, maybe precise, but may not be right. I think that's part of the conversation that we tend to have when you mix sociologists like myself alongside technologists who are trying to get to a goal.

I want to say for the purposes of our work at Brookings, we look at bias in terms of definition as two similarly situated people, places, and objects receive either differential treatment or disparate impact. And I think most of us would suggest that we look at the nature of the technical cadence behind algorithm, there's always going to be differential treatment because the models are designed in respect to micro targeting or identifying or gleaning the specifics about ourselves. But as the presentation has suggested, it's that area of disparate impact which is why we're having this conversation today.

A disparate impact within the United States is a really interesting concept because it's actually historically laid out. It starts with the values and assumptions that developers already come to the table with in terms of the models that they're building because we don't live in a vacuum. But it also lends itself to what data are we using to make the decision, where is the context, as Adrian said, of that decision.

You cannot make a decision around criminal justice or bail or sentencing in the United States without knowing the historical consequence of the over criminalization of certain groups. It's just that simple. You cannot make a decision without employment without understanding where there's been disparate impact when it comes to employment.

So I do want to, you know, really commend my colleagues on this idea that you need process and protocol that allow you to discretely discern whether or not you are engaging in unethical or unfair treatment.

In the United States we have one thing that you all have that we don't, which is a privacy standard. I mean so that's part of our challenge. There's no fundamental conversation around what federal privacy is in the United States, we're working on it. And as a result of that, without a federal privacy standard, we're on the baseline. And so we're pretty much a wild, wild West when it comes to data collection and we don't have the appropriate regulatory entities or discussion around that now that can enforce that federal privacy standard. I want to make sure that that's really clear, that that's a big difference when it comes to the type of data collected and how it's used.

I also want to suggest, and I like this idea of civil rights compliance, and, Alex, I'm going to get to your question, I promise. That in the United States there is not this explicit connection between civil rights statutes and the technological world. And so as it was suggested in the U.K. model, we don't have a stream of this compliance with those laws, this is something I'm actually working on right now, in a way that makes sense with the digital world.

Obviously we fought lunch counter bias where there was physical invasions of peoples' civil rights. Today that looks very different on the internet. Right? And so the extents to which we can apply some of those rules or the caveats of those rules is really important. And we also have, as in the U.K., sectorial specific rules that we would also have to explore and examine.

But I would also suggest that the way that these models are actually developed, we have something else which goes beyond protected characteristic. And that is the inferential economy of data that gets fed into these models. So I always give common lay people this example of a snake at the bottom of an ocean. The snake may start out going into the water, but by the time it picks up all the gravel and the pebbles and the coral it looks completely different when it comes out. That is the nature of the integrative process Adrian was talking about when we look at algorithms.

And when you combine, for example myself, as not just a Black woman, but a Black woman who's a mother, a Black woman who loves to purchase certain things, that digital composite has a potential as researchers like Latanya, Sweeney, and others have mentioned, to create outcomes that are either unanticipated or unintended, again creating a level of differential treatment based on my race as well as these other combined attributes or disparate impact where I'm actually served credit card offerings or jobs where I don't see certain things on the Internet because of the way the model has actually profiled me.

So this high level of information surveillance is very interesting to me. And so again I'm very excited to have this because I think in the United States our focus has been on explainability and accountability versus process. And that's something that hopefully we'll get to as we share the work that we're all doing in terms of how do we solve this.

So in response, Alex, your questions about criminal justice and facial recognition. And this is the challenge. If you take what I said as carte blanche, it doesn't work, it doesn't work well. When we begin to think about ways that we tease out the obvious signs of discrimination. I love it when I sit with engineers and technologists who basically say to me, "Nicol, but Nicol, the science is actually perfected to reduce the variance of folks that otherwise in person because a judge could see them, would be incarcerated."

But I start with this statement always which is, “Yes, the science is more precise, but the politics behind it are not.” And when you start with that flawed environment, that content in which the model is deployed, you may reduce the ratio of the variance of over criminalizing Black men for example on bail and sentencing, but you are contributing to a system that is already unfair.

So it goes back to what Adrian sort of suggested, or what we’re suggesting in our work, that fairness has tradeoffs. And when we look at government sectors, they too have to determine what is the tradeoff that they’re willing to live with as well as what are the processes that they have put in place not just on the design side, but on the procurement side to ensure that those tradeoffs are not massively deep to contribute to systemic inequality.

And I think that’s where government has a role to play when it comes to sensitive areas like criminal justice, like housing, employment, education, where the consequences of the predicted decision have the ability to essentially involve in the malfeasance of existing laws that have already set people back. And that’s why I think this conversation that we’re having today is so important because algorithm bias for a very long time has sort of been put to the side and technical cadence has been prioritized. And really what we need to have a conversation around today is this interdisciplinary multi-stakeholder approach that provides the appropriate data governance in addition to processing protocol to ensure that we at least reduce the potential for the type of exploitation of citizens as a result of discriminatory algorithm models.

I’ll stop there.

MR. ENGLER: I just want to say how much I appreciate the general framing, Nicol, and I think I’m going to come back to in just a second, a part of that which I was just speaking with our guests right before the panel, that is especially in policing and criminal justice I think may be the hardest and sort of thorniest part of this problem.

To stick on these topics specifically for a second, I’m wondering of our CDI guests can tell us what algorithms in the policing and criminal justice space they found the most concerning. You mentioned heat maps and predictive policing, individual riskers, there’s the use of criminal justice algorithms. What stuck out as really potentially problematic?

MS. MacDONALD: I can take this and, Ghazi, please do jump in. So I think, as I said, the sort of use of especially machine learning algorithms in U.K. policing is not sort of as widespread as sometimes we might think if we sort of read some of the research.

But there are some interesting tools being developed around sort of predicting the likeliness of someone reoffending and also around sort of predictive modeling around sort of individual risk profiles and also to help an officer sort of triage offenders as well.

So I think the most sort of interesting ones to look closely at are ones making individual level predictions because the sort of reasons that Nicol outlined around the impact that sort of false positives or false negatives could have on individuals in terms of their sort of personal freedom but also more widely on the sort of public safety. So I think those are the ones that we would highlight as to the ones that need more close attention.

But I would say most of these are sort of under consideration rather than live deployment at the moment. And of being discussed into the one police force specifically in a sort of quite transparent way through an independent data fixed committee that has been set up to sort of scrutinize unchallenged those algorithms.

You'll have to wait, Ghazi, Adrian wants to come in now.

MR. AHAMAT: Yes, thanks, Lara. I just want to add to that one sort of general characteristic to be mindful of in algorithms in policing is being careful about the target variable. Because often when training data sets, whether it's for hotspot policing, heat maps or for individual risk assessment, you necessarily have a skewed view of the world, which is what police have done in the past. And so in my trite way of saying that is predicted policing predicts policing, not crime.

And so we need to be careful about -- and there is increasing awareness of this issue amongst the people who are building those technologies. But being mindful that often these algorithms that are claiming to identify where a risk is, often is trained based on what police have happened to look at in the past.

MR. ENGLER: Oh, Adrian, did you want to come in?

MR. WELLER: Yeah, I just want to add a few comments. One is many of these issues of course are similar on the U.K. and U.S., and we are all facing many of the similar challenges in trying to understand what different stakeholders want, what they're concerned about in trying to build tools and regulation which is going to work for everyone, which is challenging.

To make it even more challenging, some additional issues come into the mix here. So I think particularly if we're talking around citizen prediction, we should let Nicol speak about what happened in the U.S. with S.B. 10, that bill raised in California that was simultaneously trying to address several challenging topics. One is how to make risk prediction and how that might be able to be used perhaps more fairly than just asking people to put up money for cash bail, that led to trouble because it made many people concerned that it was going to force courts to use algorithmic systems, perhaps before they were ready to understand all the limitations which they had.

And just one other point I'll make is Ghazi is quite to say that predicted policing tends to predict policing rather than crime because algorithms are trained on data sets and now they'll learn to predict the sorts of patterns which they can find in that data.

But I do want to correct a perception that many people have. So because of that many people have this idea that well if you have bias in, you're going to get bias out. And I say that's not necessarily the case. So we should remember that to the extent there is bias going in, it's typically because of past human bias in their decisions. If you just let them keep doing what they've been doing, you're going to continue that bias. Whereas if you use an algorithm at least with the algorithm you can introduce methods which are being developed over the last few years and we continue to develop, to try to measure and mitigate that bias. So it's usually easier to adjust an algorithmic system to try to remove bias than it is to try to adjust human decision makers.

MS. LEE: Can I jump in, Alex? And I mean I think that that is exactly where it's groundbreaking too, and I appreciate the clarification from everybody. Because part of what we see in, and, Ghazi, I'm going to always quote you on this, I loved what you said, "Policing algorithms reflect policing, right?" And so the extent to which we actually look at these flawed systems or systems that, you know, not all as my colleague, Rashawn Ray talks about, not all cops are bad cops, they're just some bad apples. But we have a historic legacy of how systems have been deployed when it comes to criminal surveillance. And unfortunately those criminal surveillance tools, and I love the distinction that we're making, break down by citizen surveillance or community surveillance.

People have not really recognized over the years that we've had high rates of community surveillance. We started that actually with traffic algorithms and public safety algorithms. Now they have actually found themselves in places where we have to make the distinction between national security concerns, much like you saw in the United States and the use of facial recognition to help us in identifying insurrectionists to the extent to which the tech locator didn't work on a Black man in Detroit who was taped for six hours and misidentified as a result of individual surveillance. So I'm thinking we need to have another panel on this, right? Because those really break down differently. I think they have different implications.

But what Adrian said I think is really key, is it breaks down to the data set, and it also breaks down to the human decision side of it. The data set obviously is flawed because if it's based on policing it's based on historic decisions that police have made when it comes to arrests. As sociologists we already know that for the most part, for example in drug crime, that whites and Blacks have pretty much equal rates of drug consumption, but most arrests are made in Black communities. And as a result of that you see this disproportionate collection of data where mugshots have more people of color represented than people who do the same amount of criminal offense. And, you know, that to me is why you need this interdisciplinary approach that we both are talking about.

But I also think on the human side, as Adrian pointed out in California, we're seeing this across the country and the world. There has to be some type of mitigation process that involves people. And that is where you have to have conversations about fairness and, Adrian, I actually just picked this up from my dear friend Professor Michael Kearns at Penn, sometimes you've got to get policy makers to understand that algorithms can also be used to solve policy concerns. Because we often look at input/

output in the black box. We solve the input by technology, we solve the output by policy. And rarely we might be thinking of a blended or hybrid way to actually get the same end if we had the right people around the table.

So I think that's another point that is well taken, particularly when you look at criminal justice and policing because oftentimes these technologies are deployed without protocol, without process, or without attention to politics. And that's where things go wrong.

We're actually working on a project here at Brookings on that, sort of with the use of facial recognition and it's deployment to law enforcement because there are no guard rails as Alex and I talk about with our team, that actually define appropriate usage and appropriate training and procurement when it comes to these tools.

So I'll stop there.

MR. ENGLER: Yeah. I actually want to stay on the procurement side. By the way, I recently read someone called this idea the iatrogenic effects of policing, which is that the intervention is what's causing the problem, right? Iatrogenic, a medical term meaning the intervention, recreating the outcome. Which I thought was an interesting take and sort of describes the problem.

Before dispensing contracting and procurement. And I want to stick on that for a second. I feel like a little jealous listening to you from the U.K. coming in and talking about how oh, I don't think there are that many algorithms out there. It does seem like the U.S. has forged ahead. And one of the areas that that creates a problem is around the incentives of private data collection, when government seems willing to procure systems, for instance in facial recognition but also in other surveillance aspects. Do we need better standards or different standards around what the governments are willing to pay for and what questions they're willing to ask and maybe what transparency they're inclined to enforce on the private algorithms that they buy for policing and other government services.

And I'll leave that open to whoever wants it.

MR. WELLER: Lara, if you maybe want to speak about public sector transparency and then I could talk about the equality duty.

MS. MacDONALD: Yes, sounds good. Yes, I think I sort of mentioned it in the slide about public sector transparency, so we as part of the report developed this recommendation around the need for a mandatory transparency obligation on the public sector. And that's sort of helping to sort of drive good practice downstream in terms of sort of ethical procurement and of these sorts of systems across the public sector. And absolutely I think it does need to start at that earlier stage and boast standards for transparency, good data quality, and sort of an explainability and do need to be sort of established at the earlier stage by the public sector body procuring the system.

And we are seeing some examples of that happening, for example some work we're doing with police Scotland where they are exactly following that approach. But it does need to be sort of looked at more broadly. And it is extremely important to set those standards at the beginning in terms of sort of contract arrangements with the private sector.

MR. AHAMAT: Yes, I want to reinforce that definitely procurement is an important link that the government has. Many of these technologies, if they are adopted by government are, in parcel, largely to private sector firms. And government has pretty significant market power to influence that. We did make some recommendations to the Crown Commercial Service, which is the part of the U.K. government that sort of centrally manages procurement standards to look at that and what sort of expectations should government or private sector providers when procuring algorithms.

I also wanted to raise a part of U.K., a quality we all call the public sector equality duty, which I don't know how widely it's adopted in other countries. But basically all public sector bodies and people undergoing a public function have an obligation to consider the inequality impacts or discrimination impact, and effectively document the thinking that they've done in anticipating and addressing potential discrimination efforts. And that's not a delegable duty. So it's not something you can say our vendors handle it. It's something that individual public sector body needs to consider. And that means they need to ask certain things of the vendor in order to satisfy that duty.

This came up in a recent case on facial recognition technology by South Wales police, which there wasn't evidence that the use of facial recognition technology was discriminatory, but they did fail their public sector equality duty because they didn't do enough to consider the possibility that it was discriminatory.

So in the U.K. at least there's quite high obligations on considering these things when they're buying from outside vendors.

MS. LEE: Yes, if I could jump into this too on the government, I do think this is such an interesting question. I mean first and foremost I think based on what we've heard, it's sort of like what are the guideposts for transparency compliance to civil rights and compliance to the commitment to doing risk based assessments around these models.

But I also think it's important for government to think about this really fundamental question of should this be automated in the first place. What we see here in the U.S. is a lot of regulatory bodies in additional to congressional bodies, who are thinking about, you know, how does AI look in financial services, how does it look in housing? So you see this constitution committee that the federal level and you have the standard bodies like this and other SPs that are trying to figure out the best practices by procuring and getting access to proprietary algorithms and setting those thresholds.

But it goes back to, Adrian, I'm going to go back to words you said that keeps in my mind when I do my work, performance of the algorithm and the extent to which the performance, when it's taken out of the context of the lab, actually performs the same among, again, different groups. Does it produce what level of impact in terms of a risk?

And I think government has not yet gotten to the performance stage partly because they are not staffed with the type of researchers and engineers and data sites that can go in and assess the algorithm from the perspective of algorithms, you know, as a decision maker in this space. But I also think that government may not also look at the variety of tools in the tool box that can help them to assess the performance, audit, impact assessments, civil rights audits, you know. When I think about government procurement I think that tells your point, people look at the ability of the algorithms to create efficiencies, but they do not understand the unanticipated or unintended consequences, and as a result they just don't have a deployment of those tools. One, because they don't have people and the resources, but two, there's just not been a conversation around that.

And so I think again, procurement is really important because procurement without everything that we said, an understanding of the tools, the training, and the follow up, basically lands government on the same front page of newspaper as it does in the private sector companies because the increased liability is not taken into consideration and these are use cases that really do matter.

MR. WELLER: Could I just add a couple of thoughts? I think Nicol makes excellent points. Just to add to that, when we're thinking about the way of course, government should lead by example and doing the right thing and doing all the things you've mentioned, audits, impact assessments, which I imagine we might be talking about soon.

You asked the question, quite rightly, we should think whether an algorithm is the right solution. Of course we should use the right solution, we should use the right tool. We should think about what are we trying to achieve, and use whatever is going to best help achieve that.

But I'd just like to note while we're thinking about this that there is a lot of attention, quite rightly, recently on how can we try to ensure that proper process is followed for algorithmic systems and that we are very accountable and we're transparent, but all of that should apply whether or not we use an algorithm. So for any policy you should be having impact assessments, you should be checking to see if things work or not, you should be held accountable for it.

So it's not that many of these risks are true whether the system was algorithmic or not. Of course we should look at them throughout algorithmic systems and there are some additional difficulties and challenges that come with the opportunities that algorithmic systems can offer. So we should keep in mind that many of these challenges were there already.

MR. ENGLER: Great. So maybe if people have a quick answer on this, quick answer on a hard question. I'm wondering how you view the overlap of historical political institutions and the resulting world of algorithms. So, for instance, in the U.S. we have an independent system of local police forces and also state police layers and federal police, you know, federal law enforcement on top of that, and some strong political actors, right? So localities have their own views on how police and justice systems should work as well as the activists groups and police unions. And so I'm wondering, do you think of these things as separate, do we create a bunch of standards for algorithms and hope people implement them? How do we reason about this overlap of existing political environments and new world of algorithms?

Maybe I'll start with Nicol if you want to jump in.

MS. LEE: Yeah, this ties to a recent discussion I've been having around state bans on some of these technologies around, policing technologies as opposed to federal bans, right, and the extent that we should have one or the other or both.

One of the things that I think is so interesting about this, I love your question, I mean I think there is the application of the type of political persuasion that happens at local, state, and federal levels when it comes to, you know, these institutions that have been historically based in, you know, local mild practices but federal compliance oversight, right?

And we see for example in the United States that for Congress who wants to actually implement facial recognition technologies in the police, can always do it for the capital police. They have no jurisdiction over other entities and municipalities.

What I have been suggesting to people though is that the challenge we're having now, this is sort of a different answer to your question, we need to go back to this conversation we're having in terms of technical cadence, right? Because we look at these technologies within the political sense, we should, because we should be very optimistic about the power of these technologies to create greater efficiencies in our system, efficiencies that we've never been able to do with the human mind. But we need to think about how do we make them better.

And so my argument is when people tell me well this locality bans this, and this state bans this, is I get it, I understand the history of policing in our country. But that doesn't mean that we should not try to improve upon the technology to make it better so that it can be applied in a variety of contexts. By banning it are we suggesting that the commercial applications and use of these technologies are not going to seek and reap the same type of consequences?

Of course they are. Because if the technology cannot identify me as a Black woman, my heritageship, and people know I change my hair often. It doesn't matter if I'm using it to open my phone or whether or not I get misidentified on the street as a criminal. And so we need to think about as we assess the type of pressure we want to put at the various level of government, this first overarching question is what are we doing to perfect the technology so that it can be ubiquitously deployed. That's the first thing.

And then I think in terms of the choice of localities to make decisions around the extent to which they get imbedded into the structural system should actually still be there. States have choices and federalisms will always usurp those choices. The key thing is are we putting in the right overarching principles of what we want these technologies to do? And I think that in the United States is not as actively engaged in that type of conversation, I see it more in an international context among international stakeholders than I see in the U.S. when we start to think about this.

And this is a perfect opportunity, for example, for the Biden administration to take on because we're centering racial equity to have conversations around the general deployment of these technologies across the U.S. but in particular where you still maintain some liberalism in terms of application.

And third, how do we continue to perfect these technologies when they are not excluding populations because there is some blur between commercial and government practices and applications.

MR. HEALY: Go ahead, Lara.

MS. MacDONALD: I was just going to give a bit of context in terms of the U.K. sort of policing landscape. So Alex's question, so in the U.K. we also sort of follow a local approach where we have 43 police forces in England and Wales. And each sort of force has a police and crime commissioner which actually has held the police force to account. The committee I mentioned earlier, the Data Ethics Committee, is of the specific police forces run from the police and crime commissioner's office. So that's a sort of example of kind of local level accountability. And in our report we do talk about the need for a more national approach in terms of sort of guidance and consistency, sharing of lessons and sharing of sort of what does and doesn't work. And also sort of retaining that kind of context specific approach that you get from that kind of local level structure. So just a little bit of context about the U.K.

MR. AHAMAT: And I just wanted to add a little bit of sort of a more general observation that we have a range of government tools to work with these technologies, and bans is one of them. We have things like audits and certification and guidance. And many of these are the tools suffered from a skills program that you need to know quite a lot about the technology to be able to interpret an audit of it. And so I think one big challenge is how do we make sure that people who do need to make decisions about whether or not to use these technologies, have meaningful information about understanding the risk, and the sorts of information that is useful for machine learning engineer to understand the risks may not be the right sort of information to give the people who end up deciding whether or not to use these technologies.

MR. ENGLER: Thanks. I should give credit, I stole that question from an audience member. That was from Ryan Watkins, professor at George Washington University, which I realize I took his thoughts and did not credit him. But thank you, Ryan.

I want to shift our attention a little and dive deeper first. So we've been talking about how one example, in policing a high states example of how government uses AI. Maybe one other typology, the other half of the typology, to how the government provides oversight of private sector AI.

In this area we could maybe consider employment algorithms you focused on in your report, and I just happen to have released a report on the auditing of these algorithms today. And it's also important because it's a good example of where government has had a historic role in oversight of the private sector.

Adrian made a point earlier that this isn't new, that we have ongoing systemic discrimination in some of these areas, employment is one of them. There is a literature review that said discrimination in the U.S. against African Americans has not gotten better in 25 years, and it's only gotten marginally better against Hispanic Americans. And so looking at what the government's role is here, is it clear what it means to what the government's role is to make these more fair, but do we know what interventions should be taking and what steps governs sort of the path from here to routinely more fair algorithms?

And again, whoever wants to take this.

MR. AHAMAT: I think there's definitely an existing regulatory framework that government then needs to know how to enforce it in the world where the technologies might be discriminatory. And another consequence of the public sector equality duty that I mentioned earlier is that all regulators in the U.K. have an obligation to consider the right discrimination in what sectors they choose to regulate.

But that's kind of on the hard end of this, and at best you're going to be able to tackle algorithms that make structural inequality worse. It doesn't tell you what you can do about an environment which may produce unfair outcomes because the underlying environment is unfair. That's a much more open question, and I don't think discrimination law is what gets you there.

MR. ENGLER: Adrian, do you want to come in?

MR. WELLER: Thanks so much, yes. So actually speaking to Ghazi's point there I think there are reasonable differences in opinion across the public about the extent to which we are just trying to avoid discrimination versus maybe going beyond that to try to remove past inequities. And those are really important discussions to have, and they're of course somewhat political.

But I just wanted to note one different aspect to this, which is that whatever it is we decide that we want to do as a society, I don't think we know the right way to try to incorporate that effectively into algorithms systems so that its ongoing technical work to trying to, if we figure out what fairness means, that's a difficult, complicated topic. But if we figure that out, we need some technical work to try to install it. But I think we also need some regulatory framework that is somewhat forgiving in enabling companies to try some different measures to try to achieve what we agree we wanted to, what we're willing to try to achieve.

So as an example, it was that famous case with Amazon where their algorithms had been trained on past human data and as a result they were biased against women. Of course that's horrible, but I don't think that the right solution to that is what I think happened, which is they stopped using the algorithms so they went back to the humans who had been demonstrated to be biased. What I think we should try to help companies come forward and not get their wrist slapped for trying to do something which is going to be helpful. We need to try to find some way to set things up, perhaps through sandboxes or some other kinds of measures that enable companies to experiment to try to do the right thing so together we can figure out a best way forward.

MS. LEE: I think on the -- so it's very interesting because the research on employment bias has been pretty settled to what everyone has said historically. I mean we've known as sociologists since the early 80s that the sounding of a person's name or the address that's a proxy actually led to the type of discrimination that people would experience. And you fast forward 20, 30 years later, we now have going back to this whole thing of civil rights protected attributes, not just the ability to see a person's face or to scan a face, but to also do social media surveillance and other types of surveillance that can actually get you back to the same end. And I think, you know, it's very challenging because these days of being able to have these blind interviews are kind of gone if we enable employers to use these other tools that are accessible to them. And we are not talking about it in this panel, but even go to the point of deep learning engagements where you're able to look at an interview and by their eye contact as well as their tonality, make decisions on the suitability for employment.

With that being said, I do think though that there is room for the type of compliance that we have traditionally seen analyzed, I think it's in your paper. And I do like the fact that we do have government entities that are responsible for providing the type of civil rights protections and compliance that need to understand this new technology. And so the fact that employment the EEOC may not be ahead of the game when it comes to the innovation of the deployment of these innovations means that we have to go back to the drawing board. And last year under the Trump administration, the administration made a big push in terms of integrating AI. This year this new administration talked about integrating AI around civil rights compliance, particularly for those entities that have the responsibility of ensuring fair and equal work places or how things access to financial services, etc.

But I also think, this goes back to, and this is where I'd love a response from my colleagues here, the training that has to be done of developers. There are companies today that work in employment AI whose developers know what the employment laws are and they have a better sense of how to build models that are much more inclusive of what those potential tradeoffs are when you violate those laws.

So I think going back to the table, being requirement of ethical or, you know, "fairness training" or an understanding of the objective variables that go into anti-discrimination is what's needed, particularly when governments are deploying these technologies, their people need to know and ask the questions of how will these break the rules. We've seen that. We've seen that with Facebook and housing algorithms that the developer allowed people to check things off as opposed to check people in, and that was a violation of housing laws here in the United States. So I would say that.

And then third I would just say the regulatory standby, and I'm a big fan of those. I think that when we get to these case where the consequences are so high, and we're not able to do the appropriate litmus test as we see for example financial services in my research, been talking to a lot of companies around, what's your litmus test for bias. And some of them do, like quantitative models in banking have a litmus test of the proxies used that suggested this is going to be disparate impact, it's removed from the model. We need to also have these areas where we're able to either over stratify our samples, ask different questions in ways that can indemnify the relationship between the private sector and government to make better models. You cannot de-bias models without having the subject that you try to impact oftentimes in those models.

And so I'm still working through what that looks like in terms of the protection that we would give people who would be vulnerable in these cases. But I do think to your point, you know, we've got to start with these anti-discrimination laws, I have to be honest. Because these are where they matter the most and they've been previously litigated to solve these problems. Because to this point we're not going to solve employment bias, we're not, in algorithms. We're going to solve it in people. But we have some tools that have already been settled that can be applied to models.

MR. WELLER: I just want to pick up on that in the recruitment context that this is one place where the U.S. is potentially at an advantage in the development of these technologies in that EEOC's sets of regulations and guidelines are a lot more precise in a way that the regulations, the laws and regulations around discrimination in the U.K. are more principles based and more context specific. There's a tradeoff there because in order to be something that you can turn into an unambiguous test that applies in all consequences, which have now been fairly well established in recruitment, about I think like the Four-Fifths rule and the like, you risk mixing on both sides context specific cases where smaller discrimination might be really important or larger disparate impact might not really be an issue.

And this is one of the tradeoffs that regulation of these technologies face. We might want these things to be unambiguous and that might be really helpful. There is a price in terms of being sensitive to context. And it's just a challenge that we need to face when these technologies get rolled out.

MR. HEALY: That said, that's really an excellent point. So I want to focus maybe on that for a second. What does it mean to be fair? There is a growing amount of research in the outgrowth fairness community around specific metrics, demographic parry and recall parry for individuals with all these numbers that you can calculate. Is there a way for these algorithms where we're going to be able to set that a specific metric where you have to hit this, the specific number, in this case will we say okay, that's fair enough? Or will evaluations of employment algorithms be more, as you said, sort of principles based or context driven where we're going to look across a series of different metrics and a series of different outcomes, maybe from the vendor itself and the company doing the hiring and say well, on that this shakes out that you're doing okay, you're doing what you can. Or is it going to be more like well you didn't meet these specific requirements enough, now there's going to be some over setter action. What's the sort of better way to think about that?

MR. WELLER: I'd suggest that first it will be useful to be able to do some kind of sanity check of different models to see how do they perform with respect to certain fairness metrics. As many of us now know, there are many different fairness metrics, over 20 that people have proposed. And interestingly, they're not compatible with each other. So you can't satisfy typically more than one of them without crushing the performance of your algorithm. But I think it's still very good from an audit or assurance prospective to be able to see how does an algorithm perform on some data tests. So I think it's a useful tool.

We should also recognize that often human decisions rely on a general human understanding of the context. And when that is the case, algorithms can't do that. At least they have not been able to do it for some good time, so it's a grand challenge in AI but we can't do it yet. So we need to recognize and embrace the fact that human and machines are better at different things. And we should be aware of the strengths and limitations of those systems and find good ways to work together.

MR. AHAMAT: I just want to pick up on that. I agree that the answer is likely to be a combination of the two. I think clarity of common definitions of how we might measure fairness and when they may be more suitable to a particular context. Like are you predicting something on a pass-fail basis versus do you care about relative performance actually changes which fairness metric might be most suitable. So I think getting common understanding of the measures of performance is going to be required. How much is good enough in every context is likely to be different. I think in some cases the stakes are so high that you really need to care about a little bit of disparate impact. In some cases the world in which you're working in is so noisy that you might tolerate by chance disparate impact. So I'm not sure we're going to get to a uniform set of standards, but we should at least be trying to aim for a common understanding of how to measure fairness.

I think the ML research community has made good progress on this. But I think there's been a translation exercise of connecting that to common sense or legal or context specific concepts of fairness.

MS. MacDONALD: Just to add a point here around sort of a more non-technical point. And I think where we've seen some backlash against algorithms in the U.K. is really around sort of the lack of public confidence in the use of algorithms in specific areas because they are so high stakes.

So I think an important thing here is around sort of the importance of securing public confidence at an early stage in the development of an algorithm before it is deployed because inevitably then you will think it is unfair that it is being deployed to make decisions about them. So that's another point to consider.

MS. LEE: Yeah. And I would actually echo what everybody said, and I'll sort of share what I've been working on in terms of these algorithmic energy star ratings. So I was really intrigued by, you know, this whole idea of SPA specification in the U.S. But more simply I was out buying a dishwasher a couple years ago, and I tell this story so people who know me know I tell it often. And there was this big yellow sign on the dishwashers that I was looking at in a big box store that told me how much water this machine consumed in terms of, you know, power, water, its efficiencies. And it had me thinking about the fact that we don't have a similar type of rating system in many respects in our algorithmic economy. Because if you go back to performance, many of the conversations that we have around whether or not fairness is actually being imbedded into the design, we don't really receive civil society's feedback, consumer feedback on whether the algorithm actually is performing the way that it said it's going to perform. And since we actually live in a more digitized economy, it seems to me only fair for people to weigh in on this conversation so I would suggest this model that I've been thinking about and, Alex, you're aware of it, which is really this three-pillar format which looks at, you know, what is the self-regulatory best practices, where is industry in many respects coming to the table with some practices even in government algorithms or government designed algorithms to suggest that they're doing primary testing, secondary testing, tertiary testing, that they're looking at different contexts, that they're making the right decisions about automation, that there's diversity in design in inclusion and responsibility that are part of the metrics in terms of the framing of the problem. Where are we looking at those verticals, those industries, domestic or internationally, that are taking the time to actually do that.

It's been interesting, and I'm sure Adrian and others can attest that through Lara, I get so many calls from companies who are putting together a fairness frameworks, who are trying to figure out, Dr. Lee, can you help us with our fairness framework? And so we're beginning to see this push towards that.

But I think the second pillar of my model is that where are the areas for public policy prescription? Where do we need government to come in and fill the gap stops, whether it's nondiscrimination, accountability, the use of, you know, algorithms in government, where there's increased liabilities around that. That to me is also important.

And the third pillar that I've been looking at is civil society engagement. So where do we check back with people. And what if we're not, when we know we're not actually optimized with the algorithm that's performing certain context, do we disclose that?

And I give this great example where people who think that their credit score is pretty solid on Credit Karma versus Experian or TransUnion, don't understand that that score is being designed by an algorithm and not necessarily by what has been adverse categories that have been litigated by law.

So to me, Alex, to your point, given the fact that fairness is so allusive, I think we need to have a framework that is adaptive enough but also at the same time protective of consumers that we can go back and revisit where the algorithm went wrong or where the consequence went wrong, have the right discussion.

MR. AHAMAT: I just want to pick up briefly, Nicol, that you're absolutely right, that industry standards and industry best practices will be a really important part of the solution. I mean government obviously has a role to play in setting the minimum standards, but in a lot of this people will want to be able to set a much higher standard. And that's going to be largely informed by the people who are closest to developing these technologies. But it shouldn't just be made by them.

MS. LEE: Yeah, I would just say I think we're also not suggesting that innovation get solved because of government's lack of understanding of these technologies, right? Where there are drawbacks or blind spots government has to be aware that they could potentially happen. But they have to have people who can go back and help them understand where that occurred.

I'm always reminded of this algorithm that was used in Allegany County that was designed to help non-physical child abuse cases. And they had a hot line of 10,000 to 20,000 people a day. They went down to 7,500 people a day. But what they missed out on, and this is similar to all the examples that we all hear about, Black kids were being kicked out of the system or being placed into foster care, excuse me, at a higher rate because the algorithm had a glitch that government did not realize was in this model. So I remember publishing that. And the developer called me and said wait, we've fixed it. And people don't know that we actually fixed this because they had a great relationship with the supplier, between government and the supplier. It's that type of nuance that I think will get us to the higher performance levels that Adrian and Ghazi and Lara have been talking about. How do we make sure that these work in a variety of concepts?

MR. ENGLER: So we are running up on time. I do want to throw one last lightening round question in from the audience that I thought was a really good one. Maybe if you have two sentences you can throw at this question that would be perfect. Is there a role for opt out here? Can governments enable people to choose to go through a traditional process, a non-employment algorithm or to not be scored by an algorithmic credit score. Is that possible?

MR. WELLER: Actually on the GDPR it's required. For fully automated decisions you have to have, it is a GDPR right to be able to opt out from a fully automated decision.

MS. LEE: Yeah. That's all I can say, we don't have that here. The way we started is the way we end.

MR. AHAMAT: I just note though that there's quite a bit of devil in the detail there to what fully automated really means. Because if you have someone who's nominally in the loop how influential are they really, particularly if they don't intervene very often, and how do they feel about their responsibility they take on if they chose to intervene once in a while. So there's a lot more discussion to have there.

MS. LEE: And, Alex, I would actually just put a little pivot there based on what I've said. I think that there should be disclosure in some respects for these algorithms as well, right? If this algorithm will have failures, systemic system failures based on context, perhaps we should start there too so people are aware how the algorithm is going to perform. And that's what I love about some of the facial recognition stuff that's coming out, this does not apply if your lighting is shady or this is the context. Letting people know, that goes back to that transparency conversation.

MR. HEALY: Yeah, absolutely. No, it certainly seems like even before opt out, just people know they're involved in algorithmic system process is something that we haven't fully addressed.

Well unfortunately that's all the time we have. Thank you so much, Ghazi, Lara, and Adrian for joining us to talk about this. For those of you who are watching and are interested, you can follow Nicol and I and many other peoples' current work through the AI and Emerging Technology Initiative that has its own little web page. I also strongly encourage you all to follow the Centre for Data Ethics and Innovation and their work and their recommendations for the U.K.

But again, thank you to our panelists. If you, I don't know, we can't see you, but if you want to like clap in front of your screen we'll sort of feel that energy anyway. And otherwise thank you so much for joining us and we hope to see you at future events.

MR. WELLER: Thanks, Alex.

* * * * *

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, 2024