THE BROOKINGS INSTITUTION

WEBINAR

KEEPING YOUR WORKPLACE SAFE
FROM ARTIFICIAL INTELLIGENCE AND SURVEILLANCE

Washington, D.C.

Wednesday, April 14, 2021

PARTICIPANTS:

Moderator:

DARRELL M. WEST
Vice President and Director, Governance Studies
The Brookings Institution

Panelists:

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

KIMBERLY A. HOUSER
Assistant Clinical Profession of Business Law and the Law of Emerging Technologies
University of North Texas
Visiting Scholar, Ostrom Workshop
Kelley School of Business, Indiana University

* * * * *

P R O C E E D I N G S

MR. WEST: Good afternoon. I'm Darrell West, vice president of Governance Studies at The Brookings Institution. And I'm pleased to welcome you to our webinar "Keeping Your Workplace Safe from AI and Surveillance."

So we live in a time of rapid change. In the last year we have seen COVID compress what might have been five years of digital transformation into five months or perhaps even five weeks. AI's been deployed in many different sectors, including the employment area. Organizations are using algorithms to screen resumes, assess job candidates, and monitor employees. There are concerns about the possible impact in terms of personal privacy, bias, and discrimination.

To help us understand these issues we're delight to have two distinguished experts with us. Alex Engler is the David Rubenstein fellow in Governance Studies at Brookings. He writes about AI deployment in the employment area and issues that we should watch.

Kimberly Houser is an assistant clinical professor of Business Law and the Law of Emerging Technologies at the University of North Texas.

If you have questions for our panelist you can email them to us at events@brookings.edu. That's events@brookings.edu. Or Tweet @BrookingsGov by using #AIGovernance.

So I want to start with Alex. So you wrote a Brookings paper recently that people can find at Brookings.edu. The paper is entitled "Auditing Employment Algorithms for Discrimination." And you say that algorithmic hiring systems are proliferating. How are companies using AI in the employment area?

MR. ENGLER: Yeah, thanks, Darrell, and thanks for hosting us today.

So there was a recent study by Mercer that said 55% of human resource leaders are now using some form of artificial intelligence or predictive algorithms in hiring. And they're using these at almost every stage you can think of in the hiring process.

This includes stuff like LinkedIn and Monster, and Indeed, which connect candidate profiles with specific job postings. It then includes a step where algorithms analyze resumes, maybe just to see probable qualifications, maybe attached to a specific job or type of job.

There are specialized assessments where you might play some sort of game that was designed, you know, it was a specialized questionnaire to see if someone will do a particular long-term job, that's another brand of company.

And then there's a whole series of automated interviews which may use things like facial analysis or may use things like the tone and cadence of your voice, as well as also the content of what you're saying through transcription, taking what you say and turning it into text and then analyzing with natural language process and that type.

So if you can think of something in the job hiring funnel there's an algorithm that's at least attempting to play a part in it. And this is where you might really start to get concerned. There's so many different algorithms playing a part in this process that if there are even small issues, much less large ones, they might proliferate throughout and build on one another and this might lead to concerns, all sorts of things, transparencies, people know what they are being evaluated on, and fairness and discrimination as well as privacy. And this is where the concern is coming from is, you know, the larger proliferation of this, the large numbers and uses of these systems.

MR. WEST: Great overview of how companies are starting to use these kinds of tools. And you started to get into this, but are there things that particularly worry you? I mean you mention issues of privacy, issues of discrimination. Like how should we be thinking about those kinds of allocations?

MR. ENGLER: The one that really sticks out to me, and this is what I wrote the paper about, is discrimination in hiring. The reason this sticks out is because, you know, I think broadly people agree that equitable economic growth and equitable opportunities in the labor market is really important. And this hasn't been going well. So a really impactful MEDA analysis in 2017 looked at 28 studies and showed that there was no change in the level of hiring discrimination against African Americans over 25 years. The level of hiring discrimination against Latinos had decreased only modestly. So 25 years with very little progress for African Americans and Latinos in hiring.

So this is to say that people aren't necessarily doing a great job. But, you know, when we look at these algorithms they're built on historical data, they're built on data from the past, they're built on top of this system that we have very clear evidence is fundamentally flawed.

And when you look at the individual types of algorithms that they're using, there's not obvious evidence that they're inherently going to be better. So we've got a flawed previous system and then building on top of this we have new algorithm processes that are likely to reflect those past biases.

We've seen this in bias against women, with people with disabilities, in resumes. We've seen this against African Americans and dualistic variations of speech. We've seen this against darker skin color, and very, very likely people with disabilities when it comes to facial analysis for young women in STEM job advertisements.

I mean you can just go on forever the number of examples. And studies is really compelling that it is very likely a systemic problem. Discrimination is the one that sticks out to me, it is not the only type of concern with the system.

MR. WEST: So, Kim, Alex just mentioned that humans don't do such a great job either in comparing humans versus the technology. And you actually have done research on bias and algorithms, and looked in particular at how the algorithms might actually be helpful in reducing human unconscious bias.

So can you tell us a little bit about your research and what you found? Actually you are mute, could you unmute yourself, please?

MS. HOUSER: I came about this looking at the tech industry specifically because they are known for having such low levels of women and underrepresented minority employees, especially in the tech positions and at the leadership level.

And I was curious to know, you know, there were some highly publicized very overt cases of harassment and discrimination, but it just seemed to me like this was more of a systemic problem and I wanted to kind of explore that and see if I could, you know, really figure out what was going on. And I had read a number of reports about unconscious bias and how that might be playing into it.

And, you know, what I discovered, there was a survey done by Human Capital Institute, and they indicated that about 80% of leaders were using gut feelings to make employment decisions. And that just kind of cemented this idea that what I think is going on is that people in general are not intentionally discriminating against women or certain groups from the tech industry, but because they have these biases of which they're not aware, they keep choosing the same type of people over and over.

And a lot of that has to do with things like affinity bias, you know. Oh, you play lacrosse, oh, wonderful. I went there too. It's just having a connection with the people who are interviewing them so you kind of develop a preference for them.

That's not to say that discrimination doesn't exist, it certainly does. But I wanted to explore more that aspect. And I have done a fair amount of legal research at a lot of these law suits and discovered that there really isn't a legal remedy for this. And, you know, a lot of people say oh, it's just for an impact, you can easily prove that. Well that's really not the case.

Our employment law is very insufficient to cover these types of subjective decision making that leads to discriminatory results.

So I looked into this idea of artificial intelligence because a lot of scholars were indicating, you know, this is really dangerous, we shouldn't be using it in these situations. And what I found was that there were companies that were actually developing artificial intelligence to actually make these employment decisions more equitable.

And they started out very big. It's evolved but they start out very basically by saying okay, let's take protected factors out and make decisions. But then they discovered that the algorithms could develop proxies and still be discriminatory. But then what they've started doing, which is kind of what Alex' paper gets into, is developing AI that could actually audit AI decisions. And that's when I really got interested because I think there is a lot of opportunity to develop the field to make employment decisions more objective. It's just AI has to be developed responsibly.

MR. WEST: Okay. Those are all great points and we certainly do need to build responsible AI. And, Kim, you mentioned some possible technical fixes in the work that you've done in terms of historic reviews, looking at the data, opening the black box of these algorithms.

Could you talk a little bit about those things and how they might help fix some of the issues that have popped up with these AI applications?

MS. HOUSER: Yes. The two main issues, and again, I over simplify it, but the two main issues that Alex alluded to are first the data. You know, we're taking this data, it's not clean, it has historical prejudices in it. And then we're running algorithms on it and surprised when it turns out discriminatory results.

So one of the things that I read a lot of papers, scientific papers, and discover that there are really ways to not only audit the databases to locate these, I guess you'd call them discriminatory databases, but a lot of times what happens is you have unbalanced data.

And, you know, the example everyone uses is Amazon. They say well Amazon developed an algorithm and it excluded women's resumes. And they use that as an example of why you should not use artificial intelligence in, you know, resume sorting, which is what they were doing.

But really it was more a case of very unbalanced data. Meaning that they were taking historical resumes which came from 90% male, using those as evidence that these are the people we hire, and then, you know, it just learned to exclude female resumes.

Well had they balanced the data either by reduction or boosting, they might have had a different result. Reduction is where you would say okay, we have, and I'm, you know, reducing the numbers to make this easier. We have 10 resumes from women, we have 90 resumes from men. You would reduce the male resumes to 10 so it's equal with the 10 for female. Or you'd create synthetic data and increase the 10 female resumes to 90 so you're comparing apples to apples.

With the black box, that's a little bit different. There's been a lot of calls for explainability, which is very difficult when you're using machine learning because you don't really know why the algorithm produced a result. But there are some technical fixes that can work on this.

One is called counterfactual testing. So with counterfactual testing what you're doing is you take your model that you've used, you use the protected factors, you find out what the output is, and then you remove the protected factors, see what the output is. And if they're the same you have a good sense that it's a relatively fair algorithm.

Another way is called quantitate input influence. And what this does, which I think it's really interesting. It doesn't open the black box but what the program is able to do is determine which categories, it lists the categories in order of importance, and then instead of, you know, having an explanation for why it made the decision, you could see what it was looking at most often to kind of guess what it was taking into account.

So if for example it was taking gender into account as number one, you know that that might not be a good algorithm to use.

MR. WEST: Okay. Alex, Kim mentioned the idea of doing audits of one sort or another. And in your paper you actually do call for algorithmic audits. But you suggest there are a variety of things we need to check to make sure those audits actually are effective.

Could you explain your view on algorithmic audits and what would be required for them to help make these tools less discriminatory and less biased?

MR. ENGLER: Yeah. This is where things, I think, get quite tricky and you have to start blending the technical and the sociological and the incentives to really get anywhere meaningful.

First, you know, at a really broad level it is true that algorithms have the potential to make strides in combatting discrimination. That possibility definitely exists. If every single hiring algorithm was really, really robustly tested with lots of investment into the process for making it less discriminatory, there are ways, you know, including balancing and counterfactual testing, though both of those have limitations, including, you know, working to improve the modeling process itself.

The data is one source of bias, but the modeling process as well can be one. You can make those things better, at least if not perfect, you can make them better. It's not obvious to me that the sociotechnical incentives are there, right? It takes engineering time of quality, well qualified, expensive data scientists to do this work. And that is work they're doing instead of developing a model for a different job or instead of developing a report to pitch a different client or instead of developing a new feature for your product.

And so while it's true these things can happen. By the way also data collection to balance, to create more representative and thorough datasets also can be very expensive and time consuming. And so you have to wonder if you look at a relatively competitive environment with a lot of new companies, are they putting in the time to make these systems definitely more valuable in terms of, you know, more fair in terms of discrimination.

They all make strong claims, all of their websites advertise on the fact that they are objective, right, they put a lot of blind faith in the idea of algorithmic objectivity, or at least that's part of their sales pitch. And that's not to say they're all the same but they are all making that claim.

And it's, you know, it's not very obvious to an employer, it's not going to be very obvious to most employers who are contracting with these vendors, which ones are really taking it seriously and

which ones aren't.

And so the incentives make you kind of think well there's probably some variation in how thorough these firms are in sort of self-assessing. And this might lead you, as it lead me, to think well what we really need to hear is some sort of independent verifications, sort of independent check on the fairness of these systems. Which, you know, just as a reminder, this is, you know, not just something I would like to happen, this is current law, it's Title VII of the Civil Rights Act, which as well as the Americans with Disabilities Act and the Age and Employment Act, right. These are various pieces of federal legislation that say you can't discriminate against some of these sorts of publications in hiring.

And we're currently just not enforcing that in any way on algorithms. And so one thing you might do is to say we're going to require, as New York City is considering doing, some sort of independent audit in which a third party comes in and evaluates on a bunch of technical levels and maybe also process levels, whether or not the algorithms being developed are fair. And that's the general idea but, yes, hypothetically these can be more fair, but are they, and how do we punish the actors that aren't putting in the time and energy to make fair, and frankly legal, systems, and reward the companies that are being responsible actors in this space.

And so far we don't have a clear incentive, we don't have a clear process for doing this.

MR. WEST: Kim, I know you've thought about the privacy issue in regard to some of these workplace applications. And we are starting to see some new policies and new regulations come on line, so obviously the European Union has taken a pretty strong stance in this area. The State of California, Virginia, have passed laws in this area.

What are the promising models that you see to deal with the privacy angle on workplace issues?

MS. HOUSER: Sorry about that. I do think that there are aspects of the GDPR out of Europe that would work well in this particular area. I also, however, think there are aspects of the GDPR that can inhibit innovation because of its strictness.

But I think to a certain degree we want people to know first that their data is being collected, what the purpose of the collection is, and how it's being used. But also with whom it's being shared. And I think people don't really have a good sense that their data that they provide to, you know,

wherever they provide it, then keeps traveling outward where they have no control over it.

But to me one of the biggest issues with this is really what I would call derived data. So you provide information, you know, voluntarily, but there are many ways that, you know, the Internet or your social media, is mined that pulls in other factors. It combines it with this information and it creates these predictions about you. Predictions about your behavior. But it can also tend to reveal very personal things about you, things you wouldn't even imagine that a machine could pick up on. And I think without an overarching federal privacy law, this is going to continue to be the case. And our data is just, it's being shared but also these profiles of our personality are also being shared and I think that presents an enormous risk.

MR. WEST: Alex, what are the promising policy models that you see? Is the EU approach the way to go, is what California has done, is that likely to be promising? What are the best ideas that you've seen out there?

MR. ENGLER: So just one quick note about Illinois, which as a state has done, I won't say necessarily a ton, but more than most other people. So on AI for job interviews they now have a law that requires notification of the use of an algorithm in an automated job interview, or an asynchronous job interview, as well as some explanation of what that system is doing.

I'm probably unfortunately quite skeptical that that information is meaningful, right. It may be like it's using your face and your tone, right? So the transparency aspect of is difficult to make valuable to people. But the fact that they know at least they've been analyzed by AI is useful. And it also enables them to request that their video interviews by deleted, and the companies must comply with that I think within 30 days. And so there's certainly some valuable parts of that Illinois requirement

The New York City legislation does go further in terms of discrimination in that it would require these biased audits. I think the really big question is how is that going to be enforced, you know. How do you require the organization that's doing it doesn't have a financial dependency and is independent and still gets what you need, which is data and model and code access, right. So a third party has to come in and really get to the internals about how a company operates.

And that's a concern to these companies. One, because if they're doing anything wrong it's going to become pretty obvious pretty quickly. And, two, because they may be concerned about the

proprietary nature of the systems that they're using.

But that's really what you need. You cannot ask questions, you cannot do this through a Q&A, you cannot just talk to the developer, you have to do a really rigorous audit. And frankly, it needs to be somewhat adversarial. You need to assume that the company is not necessarily complying perfectly or not necessarily giving you all the information.

And that's why I am very sympathetic to calls in the U.S. for the EEOC to get involved. The EEOC, which is in charge of enforcing at the federal level employment discrimination. Currently, you know, there was a letter sent by a number of senators to the EEOC saying hey, what's preventing you from going and doing a big data collection from these vendors of algorithmic hiring tools and seeing what they're doing and maybe if there are problems here? And the response letter is not public, but from my understanding the EEOC says they're not really sure they can do that, they're not really sure that they have the capacity to do it, they don't have like the actual staff talent to do it. And that they usually respond to individual complaints about discrimination.

Now that's kind of a problem here. One, because individual complaints about discrimination are hard to prove at all, unfortunately. It's very hard to take these to court as an individual even if you're sitting face to face with someone and prove racial or gender discrimination successfully in a court, even though we obviously know it happens.

But you as an individual, you don't even necessarily know, in most states outside of Illinois in the U.S., whether or not you're getting any, you know, whether or not you're being analyzed by an AI, whether it's analyzing your voice or your face or your language or your resume. And so you may not even know it's happening and thus it becomes very hard to prove discrimination or enough to go to the EEOC and lead to an investigation.

So the individuals sort of have lost influence to start the process and the EEOC, which I think personally should be doing more on this, seems like they're waiting for something that can't necessarily happen. It's not exactly clear what breaks that logjam, hopefully it may be Congress saying well, okay, here's some funding and here's a mandate, but it's time that someone went and did that. In the United States that's what I'm most hopeful for, which is a role of the EEOC and sort of an active investigative role.

MR. WEST: And that is a great point in terms of the agency expertise and some of the problems in that area because we are starting to see federal agencies get tougher in their enforcement practices. But in some cases they don't have the actual expertise to follow up effectively. So that's an issue that policymakers really need to be thinking about.

MR. ENGLER: Can I say one more thing, one note on that, Darrell?

MR. WEST: Sure.

MR. ENGLER: There is a really interesting development led by U.S. Digital Service and the Office of Personnel Management, and I think I'm leaving out some agency, in which they jointly hired like 13 data scientists across the series of agencies, which attracted a ton of interest and a ton of applicants. I honestly think someone who has taught many data scientists who are interest in public service, there is demand for people who want to go to the federal government and do this job.

The problem, in my entire experience working in this field, has been the government demand, the ability to hire people with the right salaries and with mentorship with the right titles. And there's some recognition already of changes coming both in the federal hiring process, moving to something that's more based on subject matter expertise, which is great, and also creating titles like the title Data Scientist, which doesn't exist formally, but now there's an informal title of that in the program. That is the type of thing you would want to do, hire people who are interested in this type of work and who have a background in being skeptical and understanding the features, that's a really good start.

And I'll mention one other quick point about capacity. This is so important and it's going to really cross all of government regulation, not just employment algorithms, is that of data access. If we can't figure out this problem both data and model access, if you're sending, you know, subpoenas, administrative subpoenas or enforcement subpoenas, and you're getting PDFs or boxes of documents back or documentation or you're getting written answers, that's not going to solve these problems. We need government agencies who can request and enforce requests for data access and then the talent to investigate that data to find and deal with problems.

MR. WEST: Great. Thank you. Those are important points. So I'd like to involve the audience in this discussion. We're already starting to get some questions. And just want to remind people if you have questions you can email us at events@brookings.edu, that's events@brookings.edu.

Or you can Tweet @BrookingsGov by using #AIGovernance and ask your questions that way.

So one question that has come in is from Irina of CBS News. She wants to know, how has the pandemic effected the employer use of these technologies? Is the shift to remote work making surveillance tools more prevalent?

And I'll just give a quick answer and then happy to turn it over to either one of you to extend those comments.

Irena, the short answer to your question is yes. The shift to remote work has made the surveillance tools more prevalent because people are working remote and so almost by definition they are using digital tools to do most of their jobs, either through Zoom video conferences or other types of online platforms.

And the issue is that many organizations are starting to incorporate surveillance in their monitoring of employees. So managers have this concern people are working at home, there's not a supervisor kind of directly paying attention, and so they are being tempted, and are actually using surveillance tools for worker productivity, to monitor what workers are doing. Are they working the hours they say they're going, how they are using their time? And there are all sorts of digital tools for this type of monitoring. From keystroke lobbying that actually monitor what you are typing, what appears on your keyboard, to the use of video monitors and a number of other tools.

So this is a problem. We actually have a Tech Tank blog on this topic of workplace surveillance. We also have a tech tank podcast that deals with the subject. You can find either one of those at Brookings.edu.

I don't know if Alex or Kim, if either of you have any comments on how COVID has affected the degree which employers are using some of these new types of technologies.

MS. HOUSER: Alex, were you raising your hand?

MR. ENGLER: Oh, I was gesturing to you.

MS. HOUSER: Oh.

MR. WEST: You're both way too polite here.

MS. HOUSER: I was just going to say, yeah, you mentioned a lot of the productivity-tracking software that they're using now that they did not use before.

But another aspect that's concerning is, you know, viewing into people's homes, you know. What if you see things behind them that, you know, as an employer you don't like, you know, political beliefs, you know, perhaps, guns hanging on the wall, you know. You know, being able to peek into someone's home while they're working does concern me just because, you know, everything in your home you're used to. You don't think of it as being objectionable or concerning to anyone.

But I think there is another aspect to the privacy issue, not just with the computer tracking, but just with having a camera, you know, on someone in their own home.

MR. WEST: Yeah, that's a great point. A lot of people are working from home, using company devices, not understanding that the law allows companies to monitor what they're doing, even if they're working at home, if they are doing it on a company device.

So people should certainly be aware of that and we need to do a much better job just in terms of worker notifications of what companies are doing in this area. Alex, anything else you would like to add?

MR. ENGLER: Yeah. I mean I think first of all this is a new topic for anyone. Adam Satariano, he's a New York Times reporter, wrote up a piece "How My Boss Monitors Me While I Work From Home." In which he voluntarily downloaded one of his pieces of software and then shared it with his editor. And it's a fun read, also it's terrifying.

And this software has GPS tracking on his phone, as well as screen capture of what his computer is doing, as well as a random, every 10 minute webcam photo of him sitting in front of his computer. And it's sort of terrifying to think about this, well this must be an exception.

There's a quote that I can't get out of my head from Tommy Wier, who is the CEO of one of these companies called Enaible. And he's been talking about how the incredible proliferation of this software extends COVID. And it goes "In the next six to 12 months it will become so pervasive that it disappears." Essentially saying that it's so common that you're not going to think about it anymore in the immediate future.

And that quote's really stuck with me and I think, you know, we have very, very few restrictions on how this data can be used. Again, if you're on a work device and if it's within working hours. And I think, you know, the limits on employee surveillance are pretty endless. We're just going to

keep getting better at building little sensors and better at building monitoring and better at creating algorithms that predict something and better at the nudges that increase productivity.

And so if we want to set boundaries on what can be done, we might want to do that because the technology is not going to create them. Technology is going to keep enabling surveillance pretty much perpetually in ways that, as Tommy Weir says, will disappear, we'll become used to them because of their pervasiveness. And so it may be a time to consciously consider what boundaries are reasonable, what places to say this is too far, we don't care if it's a work computer, we don't care if it's something, you can't take pictures of people's, you know, bedrooms, right? It might be time for that type of reflection.

MR. WEST: Yeah, because you're exactly right. There are very few boundaries right now. And in a lot of cases workers don't even know what type of monitoring tools are being used, how long the information is being stored, and to what uses that information is subject.

MR. ENGLER: One quick note on that. I think there was a moment where a Microsoft set of tools created essentially a surveillance dashboard of employees. And they just added it to like Microsoft Office. I think, I could be slightly confusing which the product, but it essentially added to an existing set of products to a workplace management tool.

And so you could have gone on continually using the same software you're using and suddenly there is now a surveillance component that you had no idea was happening. And, you know, that could be a concern for, you know, any cloud-based tool or any chat thing where you think about like slack. Obviously emails are being monitored in quite a lot of places.

And some level of surveillance I think, you know, there's been a few studies that are saying 50, 60, 70% of companies surveyed are actively using some form of sort of digital surveillance. So, yeah, it's certainly pervasive.

MR. WEST: So getting back to the issue of employment algorithms --

MS. HOUSER: Let me just say one more thing about the Microsoft tool. If your company does choose to use that, as an employee you have no option to opt out.

MR. WEST: Thanks for that addition. So getting back to the use of employment algorithms, James has a question. He wants to know what advice either one of you have for job seekers

who basically are being subjected to these types of recruiting and management practices by either their potential employer or their actual employer. Are there things people can do to protect themselves?

MR. ENGLER: Yeah. Yeah, it's so tempting to say run in the opposite direction. But again, that's not an option. These things are so pervasive for many jobs you can't just avoid them as an individual.

For resumes, there are tools that will analyze your resume to see whether or not it will look good to other algorithms analyzing your resume. There may be free trials for those and they may cost a small amount of money. So as an individual you might want to try and find one of those to get an algorithmic assessment of your resume before you pass it to other algorithms.

It may do stuff as simple as say having a resume broken up into two columns is not as easily read, or something, right. Which is a completely ridiculous, there's no signal, there's no value in that. But that might be the case, that that's not as effective a communication strategy to an algorithm resume tool. That's one.

For video interviews there are some guides online for instance, you know, figuring out how to make sure you're facing a camera and making eye contact is one. I feel gross giving these suggestions, right, because it's sort of you're forced into this ridiculous scenario.

And another is to make sure you're well-lit with it. There's a relationship in some of these systems between lighting, which also by the way has a clear racial connotation in that it disadvantages people with darker skin, sometimes faces don't even appear as faces in the algorithm. And so lighting can be one.

There's also some evidence that like tone and cadence can be helpful. So speaking clearly and in a, you know, sort of reasonably consistent interpretable speed. And then also, I don't know, maybe trying to sound excited at certain points.

But honestly a lot of this is guessing and you don't really know what it's interpreting. And I'll say that for some people you can adjust to these systems to some extent. It's a fair question and I understand why people would want to do that. But it's worth recognizing for many other people who have an atypical speech pattern or any type of atypical facial structure, as well as we already mentioned, accents and dialects, or disabilities or even scars or skin condition. There's very little you can do about

that, right? And so there are some ways you can get them but they're probably unfortunately systemic unfair, and an individual may never know, or very unlikely to know why they didn't find them compelling or found them to be a bad fit for a job.

MR. WEST: It's a commentary on our times that we have to sit here and provide how to advice to job seekers on how to deal with these algorithms.

MR. ENGLER: And you'll note none of that had anything to do with your qualifications for the job, right. Not one thing I just said had anything to do with, you know, be better prepared, right?

MR. WEST: It was all style at presentation.

MR. ENGLER: Yeah, wonder what happened to it, yeah.

MR. WEST: Never actions. Kim, any thoughts you have on this?

MS. HOUSER: Yes. With resume screening there's actually websites that show you how to cheat. And one of the more clever things that I discovered, and they must have learned it from a college student. But you take the key words out of the job ad, you put them in your resume, but you put them in white lettering so the algorithm picks up all the key words, even though, you know, if someone was looking at your Word doc they wouldn't see them because they're all in white.

Students do that to get the right page count. So at the end of their paper they'll type a bunch of nonsense with white lettering and oh, they met their word count. But, you know, you don't see it because it's all in white.

MR. WEST: Interesting. We have another question from Harout. This person wants to know, should the criteria that are being used in the employment algorithms be required to be made public?

MS. HOUSER: That's a really good question. I think to a certain extent you could do that, but with machine learning it's so difficult to know what exactly the algorithm is measuring or giving prominence to. And in that case the company would have really no ability to provide that type of information. I think it would be difficult.

MR. ENGLER: Yeah, I agree. One, because many of these are black boxes, they're using some random combination of features, some years of experience that it got from your resume and something you put in on a form about your level of interest in the type of work, but also the tone of your

voice. And these things get passed into crazy nearly random combinations and somehow come to the sense that you're qualified or not. And so could you explain that situation? No, definitely not. It's not even clear anything meaningful about it.

And you know an important thing to remember is that sometimes the system is good. This isn't to say all of them work like that. Some of the assessments, for instance, and some of the questionnaires, might actually have clear connections to job skills, right? Your ability to solve like a logical problem solving game might have some connection to your spatial awareness that it's valuable for a certain type of job.

There's some that say there's like no connection between algorithmic assessments and job outcomes. But in many of them it's going to be very, very hard to tell you what exactly happened. And it's not super clear employers care. An important thing to recognize is that these systems are fast and they require less human time, so they're often cheaper than getting -- and you can also evaluate many more candidates.

And so if you're a company that's hiring a whole ton of people, the point of this isn't necessarily like the most obvious value ad, it's not that it's fair or better, it's that it's faster and cheaper. And that's an important thing to remember. Sometimes people say oh, they must be less discriminatory, or they must effective, or why would people use them? The cheapness must be an important metric as well.

MR. WEST: Okay. David has an interesting question. He asks, if there is discrimination in hiring, why does it matter whether it's done by AI or by humans? And he points out the psychologists have found that most humans make decisions by intuition not on a rational basis, and that explanations are just rationalizations of non-rational decisions. And so he's worried that human decision making actually is no fairer or any more transparent than AI. And suggests that the important criteria is the outcome in terms of the real measure of discrimination.

So any reactions to that observation?

MS. HOUSER: Well there are some similarities with how discrimination law works in the sense that it doesn't adequately cover subjective human decision making because you're really trying to see in someone's head as opposed to disparate treatment where you clearly see the results of choosing

one group over another.

But it's kind of the same thing with discrimination laws being applied to algorithms. You know, I don't think there have been any employment lawsuits based on an algorithmic decision as of this time, and I think it would be very difficult to prove. In the first case it's very difficult to bring a prima facie case because what are you going to point to? Just because you were not selected for an interview does not provide enough evidence to bring a case to say that you were discriminated against.

And then even if you were able to somehow get past that barrier, all the employer has to do is say well, it's a legitimate business decision to use the algorithm, and that's their defense. And then the burden shifts back and you have to show that there's a better way to do it. Well if you don't have a computer science background and you have no access to their algorithm, how can you possibly show there's a better, less discriminatory method to choose people? So a lot has to be updated.

MR. ENGLER: Yeah. Yeah, that's a great point. I mean so I'll just echo one thing and add one point. Which is, you know, it's sort of possible to prove discrimination in a hiring process as an individual if you're staring someone in the face and they say something incredibly racially, or someone emails you and say well, the issue is your disability, right. And like if someone tells you, I mean that's very uncommon, but it's within the realm of reason that they can prove it.

As an individual up against an algorithm, there's functionally no chance there, or at least no one's done it so far. It's not clear how you have any inclination that it was happening. It's almost completely impossible. As an individual without access to data and models and, you know, without really any transparency in the system at all, it's just sort of a common issue. You're going to see it in credit scoring as well, where you have just no idea what's going into the model.

And so maybe you lose a little bit of the private litigation angle, maybe a little bit, but the way I more typically think about this is that this is an opportunity, right? We have a real opening to systemically reduce the amount of bias in decision making as we shift toward algorithmic approaches. I'm not arguing against the algorithmic process, I don't really care whether people use the algorithms or not. Honestly I don't think it's really the federal government's role to care if people are using the algorithms or not.

If the companies decide that they are more effective or faster or cheaper, sure, fine. But,

if they're going to, suddenly we have a very real opportunity to enforce discrimination laws in a way that has functionally never been enforced before. And I see absolutely no reason the concept, the federal government or other governments that shouldn't say, hey, yeah, if you're going to use this cheaper thing, we're going to enforce these laws that have been on the books for decades that are incredibly important for an inclusive and equitable economy.

So there is some real truth to the sense that I think it's very hard to look at the proliferation of algorithms and say things are definitely going to be worse or definitely going to be better. The public just doesn't have enough information to make that judgement. There's definitely enough evidence to suggest concern, and I think there's a very strong reason to say that an intervention by governments and to create this sort of third-party independent auditing and some accountability would be of great benefit.

That being said, yeah, by default is it better or worse? I don't know, that's sort of an unknown question.

MR. WEST: A Wanda Vanna has a question about the cybersecurity aspects of work surveillance tools, especially with people working remotely from their homes. She wants to know, can these tools be hacked to allow third parties to track employees, as opposed to companies just monitoring their own employees? How big of a problem is hacking here?

MS. HOUSER: Of course it's possible. You know we've had people be able to hack self-driving cars. I mean I can't imagine they wouldn't be able to hack into your home computer.

You know one of the problems with remote working is that at least, you know, in a lot of big companies you have some type of cybersecurity measures on the servers and on their internet, whereas at home, you know, you've got whatever service provider you're using, you don't have the same level of security on your home computer that probably your work has on their computers. So there's a much bigger risk. In fact the risk is a bad actor can get in through your home computer, find a way to get into the company's computer. To me that would be a huge risk if I was having workers work at home without sufficient security protections on their home computers.

MR. WEST: And of course a lot of people working at home over Wi-Fi networks, which are far less secure than a hardwired network you may find in your office.

Alex, your thoughts on the cybersecurity and hacking aspects of workplace surveillance tools.

MR. ENGLER: That's a good question. You have to assume that there's going to be a market for this sort of data breach, right, and into observation that includes GPS location and camera data of people who are working. It's a lot of information about someone's house.

You know, if you're on a work computer and you check your bank accounts, right, and you log in briefly into your bank and you put in your information. You know, you can imagine that there would be a very strong incentive for a cybersecurity breach, and I think past evidence has shown us that if there's a strong incentive to get at the data, someone will in some capacity.

It's also, you know, sort of in the same vein, it's very hard typically for a company that is, like you're a company that builds widgets or whatever, or maybe edits documents, right, and you hire, you use one of these things to observe people in their homes to see if they're really working, if they're on slack, if they're editing documents. And that's probably the thing you're evaluating. You're like is this a good surveillance software. The company and the individuals, they're just no chance they're going to be able to evaluate its cybersecurity performance, right?

And so this is like the entire market that probably, like many markets, doesn't enforce and doesn't reward good cybersecurity. And so what happens is these companies proliferate and the ones with good safety features like maybe they're the ones that sort of win the competition battle and maybe they're not. And so like you have a market that doesn't really always reward cybersecurity, and then a strong incentive to hack into these systems for personal information.

And so, yeah, I would assume it's a problem at some point. And I haven't seen those stories yet, but I think it's reasonable to think they'll come out in time.

MR. WEST: And of course at the beginning of COVID we all immediately thought this was going to be a four to six week thing, now it's a year and even longer than a year. And many of these remote work features are likely to become a permanent feature of the landscape. So what we initially might have thought was a short term and temporary adaptation to the pandemic could actually become a long-term feature. And so therefore we need to be thinking about these issues, loss of privacy, surveillance, cybersecurity risks and so on.

Molly has a question about whether we need a regulatory agency that can oversee, monitor, and enforce privacy and keep people safe from some of these algorithms. Now she does not draw distinction on whether we need a new regulatory agency or whether one of our current regulatory agencies could perform this task.

But either one of you, do you think we need someone to step up to oversee and monitor what is taking place with these areas?

MS. HOUSER: There have been some calls to have the FTC expand their authority. Now some argue they already have the authority given that they monitor unfair business practices, which could include what private corporations are doing with your data.

And then on the other side, of course, people are calling for a brand new agency that is completely in charge of this. But what I would want to see first really is the updating of the laws that apply to the government's use of data. The Privacy Act of 1974 was updated in 1988, and that was meant to prevent the government from collecting all this data from people. They passed it right after Watergate so there was a huge concern of government spying on people. And I think with this being written prior to, you know, widespread use of the Internet, prior to social media certainly, that there's a lot of things going on with the government collecting our data that, to me, is not the intent of the laws we have on the book.

But we just need to update all these laws regardless of whether we have an independent agency that's monitoring government and commercial use or, you know, the FTC with commercial use. I think there has to be some kind of, you know, someone has to create regulations because what we have right now is insufficient.

MR. WEST: So basically you think we need tougher policies and enforcement?

MS. HOUSER: Yeah.

MR. WEST: Alex, your thoughts on that?

MR. ENGLER: It's a really good question. You know I think the most important thing is to start by enforcing current laws that exist. Some definitely need changing at the regulatory level and the specific example. We've been talking the EEOC has a set of rules called the Uniform Guidelines which are sort of its interpretation of the, its regulatory interpretation of the laws that govern its oversight. And some of those just need to be reinterpreted in ways that reflect how algorithms work. Because you know

one of the things right now says if a value of something or a feature of a person is predictive for their ability to do a job, then it doesn't count as discriminatory.

But the problem with machine learning, because everything's predicted. It's definitely predictable. So like at the very least there's a whole bunch of stuff that needs to be updated in light of what algorithms are doing, that's like the minimum step.

I also mentioned the need for capacity, right? At some point we're going to need new people to come in, probably many more than we've hired so far, to do this type of oversight, as well as some level of infrastructure, right. Secure data access, the ability to go get this data and the model from companies, right, the literal legal authority to do that without having to go to a court case for every single new data collection you're trying to do, as well as the ability to store it securely. And it's a problem if the government collects a whole bunch of corporate data and that's hacked, that also is a very bad outcome for, you know, corporate espionage perspective and also with individual privacy perspective.

And so those things kind of have to happen. I know little agnostic to the FCC or not FCC. I think there is, you know, one argument for a new regulator, that you need a sort of new culture and a new set of skills, and if you look at like what the Consumer Financial Protection Bureau did, it was able to do a lot because it was building up sort of a fresh approach, a quantitative approach to financial analysis, and there was some value there. Some people have criticized the idea of a new agency, which is easier to create corporate capture, that is the companies that are doing this will be focused on influencing single agency, which the FCC has arguably been more resilient to.

So there's really good arguments there. I don't know if I have a really strong opinion so much as I'd like to see this enhanced capacity exist. If there's any really significant thing I feel that lots of this involves domain expertise. So if you're going to, you really have to understand the subject matter. So if you're going to improve algorithm hiring systems, you have to understand hiring systems. Which means the EEOC has to be involved. If you want to really improve how colleges are using algorithms to bring in students and to award financial aid, you're going to need to understand financial aid systems and then the Department of Education should be involved. Or you could say the same thing for electronic health records and bias in health systems and help in human services, and you could go on.

And so sort of a little big agnostic to where exactly the investigatory capacity falls, so long

as its joint with subject matter experts and it has these features that we need, the data collection, you know, legal authority, the technical capacity to secure data analysis, and then also the talent of the civil servants.

MR. WEST: Okay. Terrific point. So on that note I want to thank Kim and Alex for sharing each of your thoughts. Lots of terrific ideas there, clearly a number of things we need to worry about, but each of you have some possible solutions in dealing with these issues.

I also want to thank our audience for tuning in. And this video will be archived on the Brookings You Tube Channel, so if you want to share it with your friends or other people would be interested, please do so.

We write regularly about these topics of AI advice and AIS government questions, and we find that that writing on our tech tank blog and through our tech tank podcast, both of which are located at Brookings.edu.

So thank you very much for tuning in.

* * * * *

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