

THE BROOKINGS INSTITUTION

FALK AUDITORIUM

UNPACKING THE WHITE HOUSE BLUEPRINT FOR AN AI BILL OF RIGHTS

Washington, D.C.

Monday, December 5, 2022

OPENING REMARKS:

SORELLE FRIEDLER

Assistant Director for Data and Democracy, White House Office of Science and Technology Policy

PARTICIPANTS:

NICOL TURNER LEE (Moderator)

Senior Fellow, Governance Studies

Director, Center for Technology Innovation, The Brookings Institution

HARLAN YU

Executive Director, Upturn

ALEX ENGLER

Fellow, Governance Studies

JEROME GRECO

Digital Forensics Supervising Attorney, The Legal Aid Society

SORELLE FRIEDLER

Assistant Director for Data and Democracy, White House Office of Science and Technology Policy

* * * * *

Nicol Turner Lee [00:18:46] Good afternoon, everybody. Oh, you all can do better than that. We're live. Hello. This is my first live event back at Brookings. Thank you to everybody who actually is here in person. So good afternoon and thank you to everybody who is watching us on the livestream. Good afternoon to you as well. Or wherever you are in the country. Hello. I'm so excited about this event. I'm Dr. Nicol Turner Lee. I'm a senior fellow in Governance Studies and the director of the Technology for Innovation, of Innovation, for Innovation Center here. I'm also the lead for the algorithmic bias work that we have here at Brookings. So I recently published some work on disclosures, and some of you who have been following me know that I've been working on an Energy Star rating, of which I just published in the Oxford University Press Journal that will come out later this year.

I'm excited about these activities for two reasons. One, I have not been on this stage in about two years. I'm nervous, which is unusual because I've been a pixel for two years. So now I actually get to see what it's like to be a real, a person, right? But I'm also excited because we get to talk to the White House. Not too long ago, there was a release of an AI Bill of Rights, which is something that many of us in this audience have been waiting for. We all know that when it comes to the algorithmic economy, that it's not fair. And we see a lot of disparate impact, as well as differential treatment of various subjects, various people, particularly those that are people who look like me, who are multi-marginalized in many respects. And the conversation is long overdue on what agency do people have?

So I'm excited about this White House AI Bill of Rights for a variety of reasons. I think it's long overdue, but it provides us with a glide path on what equity looks like. And so before we jump into the panel discussion, because I'm sure some of my distinguished experts may have different opinions, I'm very interested in having this audience hear more about what the White House's intent was and what they mean by that. So we were very fortunate to have Dr. Sorelle Friedler come and join us, who develops and leads the AI and automated system policy as it relates to society, equity and ensuring democratic values are protected in the use of data driven technologies. And that includes the blueprint for an AI Bill of Rights Initiative at the White House. Dr. Friedler is an associate professor of computer science at Haverford college. Her research background focuses on the fairness and interpretability of machine learning algorithms with applications from criminal justice to material discovery. She was

previously a software engineer at Google and holds a Ph.D. in computer science from University of Maryland College Park.

So with that, we're going to hear a presentation from our distinguished guest. And then after that, we're going to have a panel discussion with some more distinguished guests to sort of unpack as the title of this event says The White House AI Bill of Rights. I'd also ask for you to stand on your hands for a minute because you have a chance to ask questions. So when we get to that point, we'll open it up for the audience to provide some Q&A. And for those of you who are watching us on live stream, please tweet those questions to events at Brookings dot edu and follow the conversation along and I'll share that hashtag when we get ready to start. So with that, let us welcome our guest, Dr. Sorelle Friedler. Come on up.

Sorelle Friedler [00:25:53] Thanks so much, Nicole. And thanks to everybody for coming and to Brookings for hosting. I appreciate it. So in October, the White House released the blueprint for an AI Bill of Rights to help guide the design, development and deployment of artificial intelligence and other automated systems so that they protect the rights of the American public. So I wanted to start by sharing some background on the blueprint. The blueprint for an AI Bill of Rights lays out five core protections everyone should be entitled to when it comes to A.I. and automated systems. Automated and AI powered technology has now touched nearly every aspect of our lives. They've brought many benefits from cancer detection to agricultural efficiency. But it seems like every day we read another study or hear from another person whose rights have been violated by these technologies. More and more, we're seeing these technologies drive real harms, harms that run counter to our core democratic values, including the fundamental right to privacy, freedom from discrimination and our basic dignity.

To answer this call, the OSTP released the blueprint for an AI Bill of Rights, which lays out five core protections from the harms of automated systems. These are, as you can see on the screen, safe and effective systems. You should be protected from unsafe or ineffective systems. The system should work. Algorithmic discrimination protections. You should not face discrimination by algorithms and systems should be used and designed in an equitable way. Data Privacy. You should be protected from abusive data practices via a built-in protections. And you should have agency over how data about you is used. Notice and explanation. You should know that an automated system is being used and understand how and why it contributes to outcomes that impact you. And finally,

human alternatives, consideration and fallback. You should be able to opt out where appropriate and have access to a person who can quickly consider and remedy problems you encounter.

The principles outline the kind of world that we should live in. Leaders across the U.S. federal government are already taking action protecting workers rights, making the financial system more accountable, and ensuring health care algorithms are nondiscriminatory. We might get more into this in the panel in a little bit, but a few key examples are already on this slide about ways that the federal government is using its existing authorities to try to take action and protect people from harm on these issues.

One of the key questions when we're thinking about governance that needs to be answered is when to which systems to apply these protections. The blueprint for an AI Bill of Rights is focused on protecting people, protecting our civil rights and democratic values, and so defines the systems in scope based on impact as opposed to the underlying technological choices made in any system, since such, such choices can and do change with the speed of technological innovation. Specifically, the framework should be applied with respect to all automated systems that have the potential to meaningfully impact individuals or communities, rights opportunities or access, defined to include civil rights, civil liberties and privacy, equal opportunities to education, housing, employment and other programs, and access to critical resources or services such as health care and government benefits. In the blueprint for an AI Bill of Rights, you'll also find a technical companion.

In the course of building this framework, we heard from people across America, from businesses to engineers to policymakers at every level, that there was a need for resources to help guide the creation of new protections. We took that feedback to heart. For each of the five core protections, the technical companion includes examples and concrete steps to build these protections into the technological design process. This includes information about why each principle is important, examples of problems that we've seen happen in practice. It includes what should be expected of automated systems; these are the building blocks that are both necessary and achievable to protect the public. And it includes examples of how these principles can move into practice; real life examples of current laws, policies and practices to drive new action. We especially hope that in our conversation today, and as you all think about policy, that you'll find that those expectations section that what useful again these are these actionable safeguards that are technologically realizable and necessary.

On the next few slides, I want to quickly go through what you can find in the full blueprint for each principle. I'm just going to take one of them as an example. So looking at algorithmic discrimination protections, you'll see that when you have the full statement of the principle, it's a bit longer than what I read earlier, and that is because it goes through in principle some of those protections that should be included to proactively protect the public. In this case, one such example is thinking about proactive equity assessments as part of the system design process. Next, for each principle, we give those examples of why the principle is important. Those are times when there have been problems in practice. So, for example, a hiring tool that learns the characteristics of existing employee pools and reproduces discriminatory hiring practices.

Next, we give that what, those expectations of how we can protect people through concrete steps of protections that can be put in place. These form a sort of checklist for action. Right. Depending on what your positionality is with respect to these systems, you can use these as a checklist for policy or technologists can use them as a checklist for part of the design process. We can think of these as those safeguards that we want to have put in place. And again, in the case of algorithmic discrimination protections, these include things like making sure that we assess our systems for demographic disparities and then do our best to mitigate any that we find. And then finally, each principle highlights ways these actions are already being taken, both across the federal government and outside of it. For example, there's an industry coalition that has put in place a sort of a scorecard to evaluate systems that are used in the hiring process for this type of discriminatory impact. And with that very quick overview, I'll close, and I believe we're heading to a panel.

Nicol Turner Lee [00:32:15] I'm going to introduce, and I think you're going last because you're going to sit next to me. Oh, you want me to sit? Okay. I know this is a little different from my pixel. All right. So let me go ahead and introduce the panel and first and foremost, let's give Sorelle a round of applause for that presentation and her graciousness for coming out. On our panel today, we have—and I'll go all the way at the end and come forward—is Alex Engler, who is a fellow in Governance Studies of the Brookings Institution, one of my colleagues, where he studies the implications of artificial intelligence and emerging data technologies on society and governance. We have Harlan Yu, my friend who was an executive director of Upturn based in Washington, D.C., who advances equity, justice and justice in the design, governance, the use of technology. Most recently, he's focused on the impact of emerging technologies in policing and the criminal justice system. And

then Jerome Greco, my new friend, who was the supervising attorney of the Legal Aid Society of New York City's digital forensics unit. They were established in 2013 in recognition of the growing use of digital evidence in the criminal legal system.

Thank you all for attending. Okay. I want to jump in and just start with and maybe Alex, we can start with you and then just come down a general reaction, right, to the AI Bill of Rights. I know you've written about it, but would love to hear because I think what we're hearing in this work, which is why I'm so excited that you're here, because this is what I think when it comes to civil rights, justice inequity, when it comes to AI, what's your reaction in terms of the timeliness of the document, Alex, to start there and in terms of the application of the, the Bill of Rights when it comes to federal agencies?

Alex Engler [00:33:55] Thanks, Nicol, and thanks Sorelle for joining us and to our other guests. So the AI Bill of Rights does come at an absolutely necessary time, and it's worth zooming out to give it a little bit of broad context. This isn't some new foray of government into A.I. policy. There is a lot of policy around algorithms. We fund it through the National Science Foundation, there's a Coordinating Council, there's a Commerce Council, the Trump administration put out two executive orders around artificial intelligence governance, and one thing that has never happened up until this is a broad contextualization of the harms and risks to civil rights and to consumers from AI. So there is real value in the creation of this document. A lot of people might be a little exhausted with A.I. ethics and A.I. principles, that is common if you work in this space. But it is fundamentally different that the federal government did it, did it so thoroughly, did it so carefully, and in a way that will drive not only governance, but also hopefully some independent use of algorithms. So it is really valuable for that reason.

One other the thing that's really useful about this is it avoids some of the missteps that other governments have made. It is sectorally and application specific. And what I mean by that is really driven by those agency actions that we just heard about from Sorelle. The fact that the Consumer Financial Protection Bureau is making steps around transparency in how you get credit scores and Health and Human Services is focused on discrimination in health care provisioning. You get better policy driven by stakeholders and with domain expertise when that happens. Right. And so there's a real advantage of the approach. There's a real advantage of this framing. That being said, there are some gaps and limitations that I think we'll probably come back to those. But I'll just start with the sort of broad it is a welcome and important step to see this from the federal government.

Nicol Turner Lee [00:35:43] Harlan.

Harlan Yu [00:35:43] Yeah, I agree with a lot of what Alex just said. I think this document lays down several important markers. In particular, you know, this document centers equity and civil rights in establishing that these are core values that have to underpin every conversation that we have around A.I., right, in how AI is built and developed. And as the document points out, whether there are situations where AI shouldn't at all be developed and used because of the harms that they might cause. I think the reason why this is so important is, you know, so much of the AI discourse oftentimes neglects or just fails to give equity and justice and civil rights issues the kind of weight and gravity that it deserves, especially when you think about the ways that these technologies harm people's material conditions on a daily basis.

I think the other marker here is that it's the White House putting political pressure and attention on technology and automated systems in ways that really underscore their impact on long standing racial justice, economic justice and social justice issues. You know, on the first day of the Biden-Harris administration, they put out an executive order on, on racial equity. And what that executive order did was one of the things it did was it directed all the federal agencies to produce plans within one year about how that agency was going to address racial discrimination.

When those plans came out earlier this year, I was frankly quite disappointed that across the entire federal agency, all the plans, you could probably count on one hand, a few fingers, really, how many of those plans addressed technology issues at all in the year 2021 when all of this is happening kind of a little bit under the surface. And what that told me was that the federal agencies aren't being attentive enough to these issues, that we really needed leadership at the highest level of government coming from the White House to draw political attention to these issues and to make it a priority across the federal government. I think it also underscored how many federal agencies were just ill equipped, right. They're not ready, they don't have the expertise. And that's something that we need to fix longer term.

So overall, I think this blueprint for an AI Bill of Rights is a step forward, but it's not the end all be all right. I think in particular, I was extremely troubled by the legal disclaimer on page two that apparently seemed to carve out law enforcement and national security activity from these principles. Obviously, in those areas, that's where a lot of these protections are most necessary. But recognizing that, you know, there are, it's the political landscape, right, that requires that kind of disclaimer in this

kind of document in order for that to come out. And that's broader than an AI issue. And, of course, that needs to be addressed. But recognizing that, I think on balance, this is still a major step forward.

I think this document in the long term will be judged not by what's on paper, but all the concrete actions that are going to flow from this document, particularly from the federal agencies. And like Alex said, you know, we're talking about prospective rulemaking, we're talking about enforcement actions, we're talking about regular, regulatory guidance and also legislative actions that need to really put these principles into practice. Right. I think one thing that in the release of these AI Bill of Rights that was, I think, not covered well enough was what Sorelle was talking about. The list, the long list of agency actions, the things that agencies were already doing and newly committed to doing in all of these specific sectors. Right. HUD made a new commitment to, promise to release guidance around tenant screening tools and how that intersects with the Fair Housing Act. Department of Education said that they were going to release recommendations on how AI needs to be used in the classroom or should be used in the classroom. And we'll see how these efforts shake out. But I think what's clear is that this document represents, you know, mile one of a long marathon. And it's really clear that the hard work is still in front of the federal agencies and in front of all of us.

Nicol Turner Lee [00:40:21] Mm hmm. Jerome.

Jerome Greco [00:40:23] Sure. It's hard act to follow both of you. You guys covered so much. And I agree with pretty much everything both you said. So when it first got released, I was, it was a little bit of finally to me, there's definitely some things I think we're going to go over later that I have issues with. But I also appreciated the amount of thought that went into it and the reaching out to stakeholders and communities and, and getting their, their feedback and their information and, and what they wanted to see and what they were concerned about, influencing what the final product was. And I appreciate that. I think a lot of times in this space that gets lost and or not even thought about at all, even when thinking about a Bill of Rights aspect of it or like an ethics issue. So I appreciated that that was not as I don't think it often is done or at least not done well enough.

At the same time, I do have concerns about how this will, will move forward. And I know we're going to talk about that throughout here. But, you know, look, I could put out an AI Bill of Rights myself and none of you would ever hear about it. And so the fact that we have this coming out of the White House means something. One, it means that there is effort, money and attention being put to this. And it also means that other people are going to pay attention to it. Whether or not that will cause

long lasting impact remains to be seen. But it has the, has the ability to do that. It forces people to acknowledge what is being said in there and either, either embrace it or have to combat it. And I think that is really important.

Otherwise, I don't think it gets the attention it deserves, especially in just the general public. I think a lot of us here probably work in these spaces or tangentially these spaces. And so maybe to us it doesn't seem as important. But I think the general public is not as educated on these issues as we need them to be, and we need them to have more say in it and be more involved and understand what is happening to them. I think a lot of times, they, they see negative, you know, reactions to the use of AI and are unaware that either it's been used or how to combat it or what to do or who to talk to. And that's, that's really concerning. And so I think this gives us a place to to progress from and it gives us kind of a grounding point to, to talk from. And so I think that's, I really appreciate having that and being able to point to that and have some credibility behind what I'm saying coming from the White House.

Nicol Turner Lee [00:43:14] So, you know, I'm going let you respond. Right. You've heard, I think, some things about the immediacy was necessary. I think you also heard that the application of these principles does something versus nothing. We'll get into some of those possible blind spots. General reaction.

Sorelle Friedler [00:43:35] I mean, I really greatly appreciate all of, all of these perspectives. And, you know, one of the things that Alex, I think you started off by sort of saying, you know, there have been a lot of AI ethics principles. Absolutely right. I'm in no way claiming that we are breaking new ground with this, with these principles and what we hope we're doing, right, you know, and as I think some of you also pointed to is putting the weight of the White House behind this. Right. Saying that, you know, civil rights and democratic values are important, and they are also important when it comes to our use of technology.

Additionally, what we also hope that we are, again, not really breaking new ground, but but adding to the conversation and helping to move the conversation forward from principles into practice. Right. We hope that the technical companion does some of that distillation work and helps provide the start, mile one, you know, a roadmap forward into what this could look like in policy, right? And into what it could look like to actually have these guardrails in place. And then, of course, as you point out, right, we are also trying to live up to that across the federal government.

Nicol Turner Lee [00:44:54] Yeah, it's a hard topic. I mean, just for transparency on the facial recognition side, I'm on the National Academies Workshop that was commissioned at the Biden EO on policing and criminal justice. So there's a report that that commission is actually doing, but it's a hard topic and Harlan, to your point, you know, I think where we see a lot of these grievances and egregious actions when it comes to AI cuts across a lot of sectors. Right.

So what I want to ask all three of you again is sort of go back down is, you know, let's go into some of these areas. Is it the fact that the White House could have done more right and just applied it to everything? I mean, I do. You all have heard me. I do like black boots. I don't want to be restricted sometimes the things that I want to purchase. But is that a Bill of Rights issue or should we find ways to do better enforcement or raise awareness? So I would love to hear from you all. Like, where do we close some of those blind spots or gaps? And again, audience, tee up your questions. If you're watching us online, tweet us at AI Bill of Rights. But let's dig into it right. In terms of where we could do more.

Alex Engler [00:45:59] I'd start to think so, there's two maybe broad categories of things. One is what could agencies have done but didn't do, and what can't two agencies do, and worth considering separately. And the good news I mentioned, it's sectorally specific, driven by the agencies. That's great for lots of reasons. It means you get expertise; you get stakeholders who are used to working with the agencies, you get better policies, targeted specific problems if the agency wants to do it, and if the agency doesn't want to do it, you get nothing. And so there is a bit of a challenge where some of this progress is uneven and driven by the perceived level of problem from that entity.

So we saw, for instance, very little on insurance. We've seen very little on educational access. I'll maybe let my colleagues talk more about this, talk about this. We've seen very, very little on law enforcement, nearly nothing on law enforcement. And those, unfortunately, advancement, there is is right now tied up in convincing agencies to do these things. So as an example, there were a few executive orders in the late Trump administration that unfortunately agencies didn't view as a priority and didn't really execute on. One was documenting the agency use of AI systems, how they themselves as a government are using algorithms and they kind of complied, but really the absolute bare minimum, unfortunately. Another executive order asked them to cover their regulatory authority over algorithmic systems. Basically, every federal agency ignored that order entirely, except for

Health and Human Services, which did a phenomenal job answering the questions. And if you want to see how complicated an agencies oversight over algorithms can be, I would encourage you to look at Health and Human Services response, because it will show you how complicated this can be for a large extent federal agency. So that's the good news and the bad news, if they want to do something, it's probably better policy, and if they don't, it doesn't really happen.

Now, there is one other part of this, which is what agencies just can't do. If you're an agency, you really want to enforce rules over algorithms. It would be useful if you could go see how they work. If you had the authority to get data and code and models about these things. If you could go to the companies that are deploying these systems, get that information and then set rules. Some agencies have authorities that let them do that, others don't. And we may need a solution going forward that evens out some of those capabilities, that lets you know, agencies that need suddenly maybe they didn't before, need some authority, some regulatory authority, some enforcement capacity on important decisions that fall into their legal oversight area, whether it's housing or health or whatever, but don't have the tools to go enforce even current law and algorithmic systems. So maybe those are my two categories, there.

Nicol Turner Lee [00:48:41] Harlan.

Harlan Yu [00:48:41] Yeah, I think just very briefly, I think the White House has a critical role here in coordinating agency actions, right? So many of these issues cuts across a number of different agencies, right? I think some of the examples that Sorelle showed on the screen, you know, like the EEOC, together with the Department of Labor, working on various aspects of how new hiring technologies are being used in the private sector. We're looking at HUD and DOJ, you know, with a settlement with Meta in May about how it's using lookalike audiences to deliver advertisements for housing to certain people and not others and shutting people out of housing opportunities. We had the DOJ and the EEOC also promulgating guidance about the use of automated hiring systems. Right.

So I think across these areas, and I think there's just so much more potential for the White House to coordinate and to encourage and to get federal agencies to really move proactively on these issues in ways that I feel like they haven't before. I think it's also up to, up to agencies to use their existing authorities in new ways. I think a lot of agencies already have the authority that they need, they just haven't applied that to our new technological world. Right. I think the FTC, for example, and the CFPB have really broad authority under their, authority to go after unfair and deceptive trade

practices and using the unfairness authority in different ways that they, than they have in the past. They have that authority, but they just haven't applied it before to issues like discrimination. And they should.

Nicol Turner Lee [00:50:21] Right. And Jerome, I want to put a little caveat on this question, because I think Harlan just triggered something for me. But I do want to ask this question, though. I mean, some of the research that I do is that our civil rights system is not actually designed to actually mitigate risk within digital discourse. And so the extent to which the White House has put out a Bill of Rights that is actually challenging the existing civil liberties. I mean, are we sort of saying that we really need to in your legal aid society, think about how these triggers actually act out in the current litigation system versus, you know, what we're trying to do when we talk about preservation of rights. Because I think when I read it, it's sort of like no one has really challenged the existing civil rights regime to maybe even suggest that what we are seeing here is not going to be applicable in the digital domain.

So Jerome, I said a lot. Right. But I want I want to push the envelope on this, right. Because you work at the Legal Aid Society and many of us in this audience don't know how these algorithms are treating us because they're very opaque. Right. We don't know if what I see is what you see and what you see. And we don't know if the current civil rights structure is agile enough to allow us to take that to court. Right. So when you in the legal aid position sort of thinking about what Alex and Harlan said, I mean, how do you answer this? How do you how do you fill that gap stop so that you can make this a much more agile and pertinent document to the federal government, and to consumers?

Jerome Greco [00:51:44] Sure.

Nicol Turner Lee [00:51:44] I know it's a hard question.

Jerome Greco [00:51:45] Yeah, it's a big question. So, you know, there's a, there's what you want the system to be and how you want it to function and then how it actually does. Right. Right. And and so often those things do not align at all. And, you know, we see just, just as, you know, give us an example. We see facial recognition being used in criminal cases. But I don't have the ability to challenge it in a case because they're not using it at trial, they're just using as part of an investigation. And since it's not being used as evidence is trial, there's no suppression hearing related to it, there's no cross-examination related to trial. And so I have actually no way of, of attacking it. And, you know, our courts, you know, with the exceptions of some judges, most of them are not tech savvy, are not

watching this presentation and trying to explain to them why something that's not coming into evidence and in the case is important for us to have a hearing on or discuss is just way over their head or they don't care, frankly, they're like, well, your client was, was chosen in a in a photo array. And I was like yes, but my client's picture, got put in that photo array based upon a faulty facial recognition system. Right. And then they went and found a bunch of people they said looked like him based upon that photo that was taken from the facial recognition system and gave it to someone else, often somebody who may be a different race.

And we know that the problems there's a lot of problems with identifications, cross-racial identifications. And that person says, oh, this is the one who looks most like the person who, who did this, right. Or I think this is that person and now somebody is arrested for it. And I could challenge the identification procedure, maybe things that come after it, but I can't get to the actual facial recognition that that started this ball in the first place. And so our courts aren't designed right now or at least are resistant to having that.

Now, if I had a statute that, that gave us the authority to do that, we absolutely could. But it doesn't currently exist. I think this you know, that is one of the flaws, as I'm sure we all know, that the Bill of Rights is not binding. Right. This is not this is not legislation. It could lead to that, and I'm hopeful that it does on many fronts, but it currently doesn't. Right. That our courts are not set up to handle these things. And in fact, they are often set up to be the opposite, to be resistant to changes, to be resistant to anything that is not something they've seen before, you know, unless it's reproductive rights. And so, you know, it's, it's difficult. Right.

And I, you know, give another example, and this is less on the AI front but still on the technology front and just showing how slow our courts are to adapt. You know, before the, the Supreme Court decided *Carpenter versus United States* saying that there's a reasonable expectation of privacy in at least seven days or more of your, of your cell tower location records or historical cell site location information records, every federal circuit court that had taken up that issue had found the opposite. Almost every state, almost every higher state court had found the opposite. And so for years, right. That, that was something, that was a pipe dream for us as on the defense side, and it took, to go to the US Supreme Court, and by the time it got there, there were so many other technologies and so many other things that had far advanced beyond that, that are now more

pressing, that we are just starting to get, to climb up through the court system and hopeful that somebody will pay attention.

Nicol Turner Lee [00:55:39] Yeah, I mean, I agree. And I think Sorelle, I want to bring you back into this. I think part of the, the efforts that we're seeing here is to raise awareness. Right. But, you know, Jerome's being kind, you know, these technologies are being deployed in already fractured, in races and discriminatory systems. So when you place tools that allow for greater precision in employment discrimination, etc., it's problematic. Right. But I do know when I think about the White House and this Bill of Rights, it's sort of also trying to change behaviors of companies too. Right. So there's a consumer protection side of it. Talk to us a little bit more about what you're trying to do, also to impress upon companies that that type of reputational risk and system technical risk is not what is acceptable in our democracy.

Sorelle Friedler [00:56:24] Yeah, absolutely. You know, part of the perspective that we're trying to take here comes out of actually standard software engineering practices. It comes out of product safety. Right. It's taking the perspective that, you know, we don't want to have to wait to the point where you are litigating this in court. You want to put protections in place before harm occurs. We want to try to prevent harm. And, you know, there's a lot of best practices in the software engineering industry that already align with a lot of the practices in the document. You know, in addition to a past as a computer science researcher, I also used to be a software engineer. Right.

You know, many of these things like testing, like ongoing monitoring, right, just making sure these systems that you're deploying work and continue to work, these are already standard industry practice. What we want to see industry doing more of is these additional practices, right, like disparity testing and mitigation, making sure that they are thinking really carefully about how humans are involved in the system. Right. You know what happens when the system goes wrong. How can you pick up the phone or otherwise get help? Right. If you want to try, try to get some recourse for a decision you think was incorrectly made about you. You know, we want to see industry pushing forward on that. And, you know, some members in industry have already been doing that. Right.

There are a lot of companies, as we talked about, the ethical AI principles earlier. Right. There are a lot of companies that have put out ethical AI principles. And in some cases, those principles have also started to change company practice. And so, you know, since we put out the the blueprint for an AI Bill of Rights, we've continued to meet with companies who are telling us more about their

practices, who are, you know, talking about something that they've learned that they could lift up for other companies to also learn how to adopt. Because we also need to be honest that, you know, some of these practices are tricky, right? They're tricky to do well. And of course, we want people to do them well. Right. And so part of this is also industry trying to make sure that they're learning how to do that.

Nicol Turner Lee [00:58:34] Now, I have this idea. It's called an Energy Star rating that I write about. For those who don't know, you know, you put on a Brookings page, a stage of scholar, they're going to talk about their stuff, even though they're the moderator. But, you know, I think what you're talking about is so right. A couple of years ago, I was speaking to a friend of mine, Jeannie Barton, and we were just talking about like a better housekeeping seal. I took that way further to try to buy a dishwasher and looked at that yellow sticker and said, you know, consumers need to know the durability of product design. Is this going to be optimized to me? If it's not, is it disclosed maybe pushing that envelope a little further could be interesting because it gives us more agency as to whether or not we find that AI to be responsible and trustworthy. Just had to put that plug in there.

Listen, Alex, where are we compared to our international counterparts, though, with this? I mean, are we a little bit behind? Is this going to keep us at equal playing field with them or is this something that we need to continue to discuss as people have said.

Alex Engler [00:59:34] So the big comparative point here worth keeping in mind is the European Union. They have a big piece of proposed legislation called the AI Act. A weird thing about the EU, it will pass. It is basically guaranteed to pass. It's sort of just imagine everything you think you know about how Congress works and think the opposite. So we're absolutely guaranteed to see this large single market, you know, EU wide legislation that governs AI. There is a big portion of it that is also focused at the same area as the AI Bill of Rights, the sort of use of algorithms in highly impactful socioeconomic decisions, they consider hiring and educational access, public services, public benefits, financial services to be included in that.

And the list looks kind of similar to when you read the AI Bill of Rights, you'll see a lot of overlap of where that concern is. Now what the EU is doing is very different. They are passing a single bill that is going to try to create new rules and standards for all of algorithms in all of those different things, as well as also their use in some regulated products like cars and boats and planes and elevators. It turns out doing that all at once is incredibly hard. So they are still struggling with

some of the core decisions about how to set such a wide path, you know, create rules for all these different applications sort of all at once, and that, these are not trivial problems and not just like they're being lazy. It's just an incredibly hard, nearly unsolvable problem.

So the advantage of our kind of mishmash approach that is a little uneven between some of these agencies is we don't have to do that. We can fit specific rules in the specific use cases, and they can be different for finance, and they can be different from commercial surveillance, which is different than judicial decision making and so on and so forth. And that actually turns out to be a big advantage. One thing we may run into down the road that could be a bit of a challenge is when our rules for AI hiring and their rules for AI hiring run into one another, which will happen especially in companies that build software that they sell in both places, but also companies like LinkedIn, which have hiring ads that work in the EU and work here in the US. So we may be on a little bit of a path, but there's also a White House and other endeavors on that, like the Trade and Technology Council, for instance, is working on some of these issues. But I'd say that's the big comparative point. Of course, there are other countries working on this as well.

Nicol Turner Lee [01:01:49] Yeah. Harlan, you want to jump in?

Harlan Yu [01:01:51] Sure. I want to double down on the importance of the sectoral approach that we're taking in the United States. Right. I want to encourage all of us to think more broadly about these issues and really take the issues that are discussed in the AI Bill of Rights out of kind of like the tech silo, or the AI silo. Right. I don't think, in my opinion, we're really ultimately talking about AI issues. Right. I think these are, for example, housing issues. Right. When a tenant gets repeatedly rejected for an apartment because a data broker is selling to every major landlord, a history where it's telling the landlord about that time during the pandemic when they lost their job and couldn't afford rent and had an eviction filing against them, even though they, the person ultimately paid, it still follows them around and it still punishes them in terms of future housing opportunities. Right. This is an employment issue, not an AI issue. Right. When somebody is looking for new career opportunities in higher paying jobs, but when they're online, they don't even see the opportunities that they could do because an online ad platform has decided that people like that person hasn't done that job very much historically and that they won't click on that ad and they won't make as much money. Right.

This is, these are health care issues, right? When you have a state that's using an automated system that determines that a disabled person is only eligible for 15 hours of home, in-home care,

when that person knows that that person needs 40 hours to survive. Right. This is, these are health care issues. And the same can be said on policing, on immigration, on education, in all sorts of areas. Right. So these aren't just kind of like AI issues. These are policy issues across the board. And I think the reason why this is so important is when we think about these issues from a specific sectoral perspective, it really opens up the door to all sorts of solutions and approaches that are just not possible when you think about it in terms of AI. Right. I think what's vital here is that, you know, we're not here just to solve the problems that new technologies have created. Right. We have these long standing racial and social and economic injustices that have been with us in our country for decades. Now technology is shining like a modern light on those issues, but we don't want to just take things back to the year 2000 before those technologies came along.

We still have all of these underlying issues. How do we use technology to get at these more fundamental issues? And I think the only way we actually do that is to look at the sectoral approach. I think the other important aspect here in using a sectoral approach is it brings an entirely different group of people to the table, right? If we're having a debate about AI ethics or AI policy, it's going to encompass a certain kind of person and a kind of community. But when we talk about this in the lens of housing or employment or health care, that's going to bring a totally different and much broader perspective in terms of who needs to be at the table, who's considered an expert to really talk about these issues.

I think it was, you know, on the day that the AI Bill of Rights launched, I think it was both symbolic and meaningful that we had Chair Burrows from the EEOC, Secretary Becerra from HHS, Director Chopra from the CFP, and Secretary Cardona from the Department of Education. Right. Because I think the, what the what the AI Bill of Rights is telling us too in the way that they've approached these issues, is that we need to put political pressure in places where I think traditionally AI advocacy has not put political pressure. And I think that puts the onus on federal agencies and also on us as a community of advocates to really expand our purview in terms of where we, where we put our attention.

Nicol Turner Lee [01:05:53] So let me just stay with you, Harlan, for just a moment, and then Sorelle maybe jump in and then I'll go to Jerome. But isn't that the intent of this, right, to get federal agencies to be more accountable, I think, are you suggesting that really, it's not just one federal agency, right. The Department of Energy could actually use algorithms to determine energy use and

whether or not they're investing resources in lower income communities versus higher income. And at the end of the day, is this an exercise for a federal government to be also accountable for this right versus putting in guidance? Because that's what I love about the EEOC. They're very clear about the punitive consequences if people use algorithms to, you know, dismay people from actually applying. Are we going in the right direction with the sectoral approach, or should we be doing more?

Harlan Yu [01:06:34] I mean, I think it's both. It's both the federal government and government use itself, especially when we're talking about law enforcement. Right. The adoption of tools in law enforcement that use AI and predictive technologies, but also for agencies that are regulatory agencies that regulate private commercial use of these technologies in housing and in employment and all the other areas that we've talked about.

Nicol Turner Lee [01:06:57] Sorelle, you want to respond to that before I go to Jerome.

Sorelle Friedler [01:07:01] Which piece of that?

Nicol Turner Lee [01:07:03] About the federal I mean, you know, really, that's what you're saying. Federal agencies have to start this, right?

Sorelle Friedler [01:07:07] Absolutely. And they are, right. You know, as you were mentioning earlier, there are lots of actions that are that are already happening across the federal government, including by enforcement agencies. Right. But I also want to make sure that we are not limiting our focus here just to the federal government, because I you know, going back to what I was saying before, right, about industry that is also critically important, and I also think exciting. Right. I think that there are a lot of chances for innovation and, you know, making this an industry competitive advantage. Right. To really make AI that people trust that respects our rights. So I just I don't want to lose track of that that piece as well.

Nicol Turner Lee [01:07:52] Mm hmm. Now, Jerome, I'm just gonna put a caveat. Now, if you have a question in a few minutes, I'm gonna be going to Q&A. We'll have some runners around passing around microphones. Please think about your questions. We've talked a lot about some parts of the AI Bill of Rights but we haven't talked about humans. And the reason you're sitting here is I had the opportunity to meet your president, CEO Twila at this AI Bill of Rights launch. And I said, the Legal Aid Society is sitting here with tech wonky people who are interested in this. Right. How do we exercise human alternatives? I mean, in your view, you talked a little bit about that in law enforcement. I'm sure there's a range of cases that you deal with at the Legal Aid Society. How do we

give this agency of the Bill of Rights to people? And where do we place humans in the loop to ensure that even as a government, we're doing this right? Because at the end of the day it goes back to community.

Jerome Greco [01:08:41] So, you know, one of the problems with AI, even with a human in the loop, is, is that human really empowered to actually do anything? And will they. Right. If you know, we have all these systems where it says, well, the, it's giving advice to the person like, like judges on sentencing and saying, you know, they could, but they don't have to follow it. But everybody knows that if you could point to, oh, I follow it because this is what the advice was, that you're never going to go wrong, even if it turns out terribly for you. However, if you disregard that advice or override it and it goes wrong, you're going to be held accountable. And people understand that. They understand that from being elected or from getting promotions or getting fired. And so we have to have somebody who actually has true, independent and will not suffer consequence when they override the system. And I'm not sure how many people are willing to do that or how many agencies would actually empower somebody to do that.

And so I think that's really a problem we run into that it's, it's almost a fallacy to say that that somebody, a human in the loop, will somehow cure the problems. And so I think that is an issue that I don't think has been resolved by this. And I don't necessarily know I have the answer for that, except for what I just previously said. And so in the criminal context and honestly, it's the most disappointing part of the, of the Bill of Rights to me was, is the exemption for law enforcement, the exemption for intelligence, when so much of the harm that's being caused are by those agencies or by those employees. And so to, to exempt them, I think, look, if you didn't exempt them, I still think it would've been hard to get them to follow it. And so maybe that was a thought in the process, why deal with the backlash if they're not going to follow it anyway? But by exempting them, I think it emboldens them because they're just so used to having everything be exempted from, to not being held accountable, not to have any sort of transparency. And by making that explicit exemption really just bolstered that.

And so I find that very troubling, especially when you think our original Bill of Rights, our constitutional Bill of Rights, the fourth, fifth, sixth, eighth Amendment were all related to controlling that power. Right. They tried to tap down on that power to give people some sort of way to fight back against it if it gets out of control. Now, there's all sorts of problems with what has happened in terms

of, you know, mitigating that ability and, you know, watering it down. But that is, I think that was my disappointment I think with, with the bill.

Nicol Turner Lee [01:11:39] Yeah. It sounds like I got to do another panel on this whole law enforcement stuff. I have to come back to this. Before we go to questions, I do have one more question for Sorelle, which is an area that I am so shocked we did not talk about, which is data privacy. This whole panel did not say the P word at all. And I need to make sure that we talk about data privacy and just making sure people understand where the White House is slipping that in because we have an unsettled debate when it comes to data privacy in this country. So I'm just curious. You put it in as well. Where do you see that as critical to every conversation we're having right now on this?

Sorelle Friedler [01:12:14] Yes, it is critical to every conversation we're having in this. You know, and that's both because of historical data privacy risks, right. But it's also because of the new risks that A.I. brings. And that's both because of the way that AI can scale, right, but also because of the way that I can lead to inferences about you, even if you are not in the dataset. Right. So we can think about social networks where, you know, even if you are not part of the social network, right, the things that, you know, your friends share about you can lead to a social network understanding something about you. We can think about ways where, you know, inferences about communities, you know, based on where you live can end up meaning that again, even if you're not really part of the data set, you are part of the data set. And so AI ends up driving a whole host of new harms that stem from this, this lack of data privacy.

And, you know, I also really there want to emphasize that one of the things that we did in scoping this document was to really emphasize the importance of communities. Right. In other words, we emphasize our concern about harms that accrue to both individuals and communities. And part of the reason that we emphasize communities there is because of this type of inference, right? This type of spreading of harm beyond just the individual who knows maybe that they are involved and, you know, being impacted by some system potentially to their whole community. And that also, I think, leads to some of the things that we can do about it. Right. Because some of the ways that we can identify these harms is by looking for them at a community scale, right. About thinking about the that full reach of data.

Nicol Turner Lee [01:13:59] Yeah. No, I think that's so important. We're about to launch here at Brookings, an anti-racist approach to A.I. and it's a pretty poignant and pioneering effort to sort of take in everything that everyone is talking about but to lead with pragmatic solution. I think I was talking to you about that. It's like it's, it's so important because we always land up here when it comes to this stuff. But data privacy is certainly key. And probably I just heard this new concept about two months ago that I've actually adopted, which is multi-marginalization. Right. When it comes to communities have intersectional marginalization and AI it makes it easy for people to identify people based on those signals.

If you all don't have questions, I do. Raise your hand if you have questions. Great. Let's start here and then we'll go to Mark. So we've got a question here. We're going to move into the Q&A part of it, and then we will begin to wrap up. I could keep you on forever, but I'm not. Okay. Okay.

Audience Member [01:14:55] My question is about giving this some teeth in Congress. So we do need some legislation and various places for it to actually be the law of the land and not just kind of aspirational on the part of agencies and other things. Does the White House have partners in the House or the Senate? Are you looking for some? I'm a Senate staffer, we could talk. This is great stuff. I hope it becomes law.

Nicol Turner Lee [01:15:16] So, you know, one of the, there are a number of different pieces of this. Right. You know, folks have been talking about the, the sectoral approach. And, you know, that may end up being true for legislation, though, of course, there's also the cross-cutting data privacy legislation that the president has been very supportive of. You know, so I think certainly that's one aspect where things are, are already hopefully moving. But, you know, there's also, there's been a whole host of bills that have been introduced around sort of various aspects of this. Right. You know, some focusing really heavily on transparency, which is, of course, key. Right. Others focusing on various sort of sector specific harms. Right. And I think that, you know, what we're going to end up needing is essentially all of these. So, you know, that's not to sort of point to any really specific legislation, but to point to the collection of them as necessary and important.

Nicol Turner Lee [01:16:16] Thank you. You want to jump in or are you good.

Alex Engler [01:16:18] Just quickly, you know, this is a real challenge for legislation. I don't know if I would expect like one law that comes around and like solves AI just the, the pure manifold of uses and all these different applications makes it challenging to solve in one sweep. You could hope

for some more capacity or some more funding for agencies to adapt. You know, cases like the Department of Education has like one you know, you can count on one hand how many people there are currently working on these issues.

So there are some capacity issues you maybe hope to solve, but I do expect it to be a long, ongoing adaptation of government to technology challenges that, as we've talked about, include algorithms, but are also within these broader policy problems. Right. So while I do think I'm hoping to see some legislation, just know that that's not going to be like we pass an AI law and we're done, right. It's going to be like a long, ongoing adaptation to technology problems.

Nicol Turner Lee [01:17:07] Yeah, there's a lot of AI laws that are still waiting to be passed.

Alex Engler [01:17:09] I don't. Yeah, I don't know if I would even root on, I wouldn't bank on any of the current ones. I'm crossing my fingers for data privacy.

Nicol Turner Lee [01:17:15] Mark over here.

Audience Member [01:17:20] Hi. Thank you, Nicole and Brookings for organizing this very timely panel. I just came back from an AI conference in Brussels and was absolutely thrilled to hear Dr. Alondra Nelson speak about the Bill of Rights on behalf of the United States. We strongly supported the initiative as you know, we were even tweeting this morning at the Trade and Technology Council, which was meeting up at the University of Maryland, urging forward motion on the Bill of Rights.

But I do want to make this point, and there is also a question. When I saw the report, it reminded me very much of a famous report from the privacy world, the 1973 AGW report, Records, Computers and the Rights of Citizens, which led almost immediately to the passage of the Privacy Act of 1974, which is still the most comprehensive privacy legislation we have in the U.S. And I think there is a real moment and a real opportunity that should not be lost, and that is to move forward aggressively on a bipartisan basis, to take the recommendations contained in the blueprint and make them into law. I think that could happen, and we have a very good precedent for precisely how this happened at a very similar moment, a new technology posing common challenges.

But my question is actually to Harlan, who I think makes a really thoughtful point about the need for a sectoral approach for AI, because we actually want representatives of the communities who are impacted by AI at the table to talk about what that looks like and what needs to be done, rather than the computer scientists telling us about GPT three and GPT four, which is, by the way,

super interesting, but it's a very different conversation. But my question to Harlan is, wouldn't you agree still that there are going to be common problems across the federal government as AI systems are deployed around such traditional questions as fairness and, and accuracy and transparency? And wouldn't you want in this moment to build the institutional capacity, maybe at OSTP, maybe at OMB, together, to push forward the type of oversight and resources and support so that as the agencies deal with these issues, they have a place to turn.

Harlan Yu [01:19:59] I mean, yes, I think it's, that kind of, you know, more resourcing, more capacity is necessary, but not sufficient. Right. That, I think that's baseline. I think, you know, we need legislation to provide more transparency to a lot of these systems, right, that give advocates, journalists and others a chance to be able to scrutinize and be able to diagnose what these problems are. But I think ultimately the solutions aren't going to come from limiting our thinking, just to technology. Right. These are broader social issues that need broader solutions. Right. One thing that we're working on at Upturn is, you know, we were thinking and talking a lot with Jerome about how law enforcement search people's cell phones using forensic devices. Right. And I think there are things that we can say in terms of legislation and regulation around what law enforcement's doing there.

But in addition to that, a lot of these searches are happening not with a warrant but based on consent. Right. And banning the use of consent searches, because there's always a power imbalance between a law enforcement officer and somebody on the street, I think should just almost never happen. But we wouldn't think about a ban on consent searches as like an AI policy issue. But I think it is, right, the same thing with the way that, you know, eviction records are being bought and sold right. By data brokers and having courts seal eviction records at the point of filing, right, such that data brokers can't ever buy that to begin with, because that actually, I think, is protective in stopping a lot of this data from flowing into tenant screening tools. And so, yes, I think we need an all of the above approach. We need more resourcing, we need more capacity at the federal agencies and also, you know, with companies as Sorelle was saying and also at the state and local government level.

Nicol Turner Lee [01:21:55] We're running out of time. Right. So I guess I have to first and foremost, say thank you to our panelists. Let's give them a round of applause. I'm looking at our events manager like, can I get another question in? But I'll be at 3:00, which is now. And I want to thank all of you for actually joining us. One thing I'd just like to do, and I like to do this at a panel is just

summarize what we heard, is that I think we're on the pathway towards some clear guidance. We needed something to sort of wake us up when it comes to A.I. and how A.I. is discriminating against certain vulnerable populations who are also multi marginalized. I think we also heard today that this is going to require a lot of stakeholders and it's going to require some self-regulatory action with federal responsibility, is I think, the balance of what we've discussed here today.

And then finally, I think an answer that we got to keep talking about is how do you get people who have the lived experiences of the communities that are supposed to be protected by the Bill of Rights in these conversations? Unfortunately, the federal government has to have people who look like the people affected. You know, that would have helped the young man or woman, particularly African-Americans and Hispanics who don't get in on an interview because they're prescreened out of the process. That right has already been violated before it becomes a Bill of Rights. So stay tuned because we're going to talk about that more in Brookings. Thank you, everybody, for coming. Thank you for those of you who watched us over live stream. And we're back.