

**Jason Green-Lowe** | [00:00](#)

All right, welcome everybody. Thank you for coming to this intense and brilliant panel discussion. No pressure. We are here to talk about the Executive Order from October on artificial intelligence. Why does it matter? Where is it going? Why doesn't it matter? What does it need in order to matter from you, from Congress?

First, a couple of words from me about... Why talk about AI at all? Why care? Who is the Center for AI Policy that's hosting this event? And then I'll let all of our other panelists introduce themselves, and then we'll dig in. We'll try and do about 20 minutes of moderated questions for the panel, and then about 20 minutes of questions from all of you, and then have some unstructured mingling so you can stay and meet anyone who might interest you.

A hundred years ago or so... We had automobiles, we had tractors, they automated away a lot of the drudge work of getting from point A to point B. It was a physical automation. About 20 years ago, we used MapQuest to print out maps. And it's on my mind right now because my phone is dead, so I'm sort of back in a 90s headspace. And the thing about the MapQuest, we almost never used Google then, was... It was just looking up information for you. You're basically just retrieving a map that helps you go where you want to go. It's automating the work of recording and retrieving information, but it's not really automating decision-making. You still have to decide how you're going and where. And today we have self-driving cars. They're not always road legal, but they're making decisions about where are they going to go, how are they going to get there. And it's not just cars either, right? We have AI making decisions about... what kind of music do you want to listen to? Or what kind of art do you want to see? Or, you know, where should I fly this plane? Or what should I stock in my store? Or who should go to prison? Or who should be released from prison? Or who should be investigated? Or where should we fire this missile? It's making a lot of important decisions... we're... automating not just the work of retrieving information, but the work of deciding what to do with that information, which is a very weighty responsibility. So Goldman Sachs thinks that there's going to be \$200 billion invested in AI by 2025, but in some ways that's... conservative, right? That's an underestimate on the impact that this is going to have on all of our futures.

I think the Executive Order recognizes how wide ranging some of those impacts are going to be because it tackles so many different topic areas. It asks, what's the effect on privacy? What's the effect on our immigration law? What's the effect on our national security? And I think that's appropriate. I think we'll need to look into all of these issues, which is what the panel is here to help us do.

At the Center for AI Policy, we're mostly focused on the risks discussed in Section 4 of the Executive Order, the dual-use risks. We're concerned that AI that's intended for peaceful and civilian purposes might be misused to create weapons of mass destruction, or possibly even cause a takeover risk, a risk that AIs could wind up in control of part or all of our infrastructure. That doesn't necessarily mean the robots with the glowing red eyes from *The Terminator*. It could mean something like social media. A lot of us spend a lot of time scrolling through social media. The algorithm decides what your eyeballs will be connected to. Well, what happens if the

algorithm also decides what you buy, who you talk to, or how you spend the rest of your time? Most of us wouldn't choose to look at exactly what the social media is showing us, but it's nevertheless wound up in effective control of the two or three or four hours a day that people spend on it. And so... At the Center for AI Policy, we're working to promote policies that might mitigate this risk, that might make it less likely that an unaccountable algorithm is making important decisions for us.

But we've got a variety of people thinking about a variety of the important risks posed by AI, and I'm very happy to share them with you. Sam, do you want to introduce yourself and we'll go down the line?

**Samuel Hammond** | [04:20](#)

I'm Sam Hammond. I'm a senior economist at the Foundation for American Innovation. You might better know us as the Lincoln Network. We changed our name back in May. My focuses at FAI are sort of on the one lane, the national security risk from AI and how they intersect with things like export controls and chips and CHIPS Act. And then on the other hand, AI for modernization and what it means for institutions, for government, for legislative branch. And that's basically what I work on.

**Daniel Colson** | [04:51](#)

Hi, my name is Daniel Colson. I'm the Executive Director at the AI Policy Institute. We're a policy research and public opinion polling group. A lot of our work is focused on trying to understand what American voters think about AI policy issues and kind of frontier tech regulation more broadly.

A couple stats from our work, just to kind of contextualize some of the conversation today. So, when asked, "do you want to go more quickly with AI technology or slow down and go slowly and deliberately?" We find 77% say they want to go slowly and deliberately, and 8% say they want to speed up the pace of AI progress, which I think generally points to the mindset that I think most Americans are in, in that they see a lot of progress happening and a lot of change in American society happening, but it's not totally clear that that's good for them. And I think a lot of people are looking to their experience with social media and the certain significant harms from that as a lesson for what happens when the tech industry oversees dramatic social changes. And so generally speaking, the American public is feeling concerned about a lot of these changes. And that's the thing that we see showing up in the polling.

On the Executive Order in particular, we found 69% supported the policies overall, and 75% said that the government should do more to regulate AI. So I think that's some amount of indication there.

Something that I think is an interesting current point of difference between public opinion and I think a lot of the conversation on AI is the American public is very supportive of straightforward capability restrictions. And even 70% say preventing AI from reaching general superhuman capabilities should be a priority of policy. So, you know, that's an interesting sort of input as the

government considers what policy should look like. The thing that we see is pretty overwhelming bipartisan support for pretty stringent policies. And that's something that I'm happy to expand on more in different areas.

**Bill Drexel** | [07:36](#)

Greetings all. My name is Bill Drexel, Associate Fellow at the Center for a New American Security. I study AI use in kind of high-impact national security domains, so your bio, your cyber, finance. Missing one. It'll come up later, I'm sure. Glad to be here.

**Elise Phillips** | [07:57](#)

Hi, everybody. My name is Elise Phillips. I'm Policy Counsel at Public Knowledge. Public Knowledge is an advocacy organization that prioritizes supporting freedom of expression, Open Internet, and access to affordable communications tools. I specialize in competition and privacy issues, and that's what I'll be talking about today.

**Jason Green-Lowe** | [08:18](#)

All right. Thank you. So lead off question, what do you see as the single most well-designed policy in the Executive Order? Not necessarily the most important, but the one that you can't really think of much to improve about it. Why is it useful? And what does Congress need to do in terms of oversight to make sure we actually get that use out of that policy?

**Elise Phillips** | [08:43](#)

Sure, absolutely. So as I said, I specialize in privacy as well. And what I think is most notable about the privacy piece of the Executive Order is that it has like a series of tiers that it prioritizes. So it prioritizes federal support. It asks that the federal government strengthen research and guidance for companies on how to comply with privacy when they're implementing these AI technologies. And it also explicitly calls on Congress to pass bipartisan privacy legislation. And so those are three things. You know, executive orders are pretty limited in what they're able to do, due to the constitutional bandwidths that they're allowed. But it addresses industry, it addresses the needs of federal government, and then it asks the legislative branch to take action. And I think that's a pretty effective way of approaching the topic.

**Bill Drexel** | [09:31](#)

Sure, I think my favorite bit of the Executive Order is how it kind of pins preliminarily advanced AI at a certain compute level, but then tasks Commerce with coming up with a more rigorous definition. I think that's helpful just because it both gives some teeth to what we're looking at right now, but also acknowledges that we are and will continue to be in a state of flux. I think they also stipulate that we should continuously revise this definition. At a time when the technology is already very powerful, but also still fast changing, I think that sort of approach is really useful.

**Jason Green-Lowe** | [10:13](#)

So what, if anything, does Congress need to do to support those updates?

**Bill Drexel** | [10:18](#)

Well, I think Congress needs to basically keep up to date with how these things are changing because it is so rapid and maybe adopt a similar approach as it's trying to enact further laws around the technology. Maybe where necessary plant a stake in the ground on some capability on some particular type of tool now, but with the provision that that can change and with particular clauses on when we should be making those reviews for that change as necessary.

**Daniel Colson** | [10:56](#)

I was actually going to make the same point regarding compute thresholds. An element I want to highlight with that is given the pace of capability changes, I think it's right that setting a simple rule in that space right now works well, but probably delegating that to an agency that can quickly change that without having to go through a congressional process will be very important.

**Samuel Hammond** | [11:26](#)

Yeah, I think you did for sort of high-level organizational purposes. You can think of the Executive Order as having a substantive component and a sort of procedural component. The procedural component is something we're used to. It's a lot of requests and reports to different agencies. So, you know, I would say over half of the Executive Order is literally just instructing agencies to explore the use of AI in their departments. You know, building some basic governance around that for procurement, for risk mitigation. There's some more substantial parts of that. So the Executive Order does a lot to sort of centralize AI within HHS, which I think is interesting because the possibilities in healthcare are interesting.

But that's mostly just interagency stuff. The stuff that's more substantial is the stuff that was just mentioned with these dual-use foundation models and defining compute thresholds. So I think this is actually the thing I was most pleasantly surprised with the Executive Order because you often hear in tech policy debates that you should regulate only on uses and not based on... be technology neutral. And that would lead you to say, well, let's not set some relatively arbitrary threshold of, like, a size of a model or something like that. Let's just focus on the uses. But the difficulty of that in the context of AI is there are companies out there whose sole mission is to build artificial general intelligence. And general intelligence, the G in general means it's not any one use. It's all uses. And that leads into dual use as well. So. One of the important things to keep in mind is—and researchers have looked at this—the scaling laws in AI, how big you scale the model, are a really reliable proxy for how capable they are. So if you recall the transition from ChatGPT-3 to -4. More or less the same architecture, more or less the same training data, basically all the text to the internet, textbooks and so forth. But they just scaled it up and suddenly it was passing the bar exam and demonstrating higher-level reasoning capabilities and so forth. And that's just a function of scale.

So the most substantive part of the Executive Order is this invocation of the Defense Production Act to require companies that are training models with 10 to the 26 flops of compute. (Floating point operations, it's the number of basically mathematical operations that go into training the model.) To report to the Department of Commerce basic safety testing and disclosures and red teaming.

There's been a lot of discussion about regulatory capture, and is this an example of that? I don't really think so, because the thing is, no company to date has trained a model that big. This is a forward-looking policy that's looking down the road to the milestone, the frontier models that will be trained in 2025, 2026. So the biggest one to date that's been released, GPT-4, is maybe 10 to the 24 FLOPs. So 10 to the 24, 10 to the 26, that's a thousand-fold difference. Regulatory capture doesn't really enter into this discussion because the whole idea of regulatory capture is "let's introduce a regulation that has a fixed cost for our competitors and then they have to pay that cost where we're bigger, we can afford the cost." Training a model of 10 to the 26 FLOPs is a \$10 billion investment. So the companies that are going to be building these models are few and far between to begin with and are going to get fewer and fewer and farther between as scaling increases. Because the more scale you throw at these models, the more money you need, the more capital, the more GPUs, the bigger data centers. And so it's really only going to be the Googles, the OpenAIs, the Anthropics, the Microsofts. The small handful of players are going to have the capital to even build these models in the first place. And all this does is add a simple disclosure or reporting requirement to say, look, this is what the safety tests of our model show before we deploy it.

Now what's controversial obviously is that the Defense Production Act is relatively untested, sort of limitless kind of law, and that's where I think we could use a congressional complement to actually codify some of this. And particularly the oversight's going to be important, because the way the Executive Order is written is this 10 to the 26 threshold is going to be a moving target. So there's a process in place for the Department of Commerce. Secretary of Commerce and a group of their peers get to update what this threshold is. They can lower the threshold, increase the threshold. They can add new variables. And so that's going to be an ongoing discussion and something that Congress needs to pay attention to.

**Jason Green-Lowe** | [16:00](#)

Thank you, Sam. So Sam mentioned these reports that take up almost half the Executive Order. For the whole panel, if you have an answer, we don't have to get every question answered by every panelist each time. But if you have a thought about a report that Congress should be excited to read—what does the Executive Order command to be written that Congress should be excited to read? Why that report? And what can Congress do with that report once it's ready?

**Bill Drexel** | [16:28](#)

I've got one. So in one of the lesser commented on areas of the Executive Order, it stipulates that USAID write a playbook for development. Which may seem trivial, but the fact is that China is taking a very aggressive approach in building out AI ecosystems across developing economies. And we're looking at a situation where we might have a Huawei 2.0 where we're trying to get people to rip and replace. Very inefficient, reinforces autocratic norms, leverages China's power across the world. At the moment, all that the Executive Order has said is we need a playbook for development for AI. So I think Congress would do very well to look at that playbook and see if we're getting anywhere close to where we would need to be to compete

with China around the world in setting, basically building out, the AI ecosystems we would need to be competitive. There are downstream effects on that as well in terms of international norms setting and all these sorts of things. But that's one area where I thought as a whole, the Executive Order was weak. And in general, we need to be strong if we're serious about countering China on these issues.

**Samuel Hammond** | [17:48](#)

One I can throw in there, the section on healthcare and HHS includes a report, a study to look at AI and drug discovery, what the implications are there. I think there's a general threat model that I have where AI, no matter where it's applied, massively increases the throughput demands on our institutions. So to give an example, a few weeks ago, Google DeepMind released a model called GNoME that predicts new molecular structures. And using this model, they basically trained it on a whole bunch of known chemical structures and predicted 2.2 million new crystals, of which like 800,000 could be useful, stable materials. And so this represented 800 years worth of material science knowledge achieved basically overnight. And now material science researchers are combing through all these candidate materials to figure out if they're technologically useful. Imagine if that happened to drug discovery, right? And we went from, you know, the FDA in a typical year approves 15 new molecules for new drugs. If that went to 500 or 5,000 or 50,000 per year... they just, they simply don't have the capacity, right?

And that's a pattern that's going to repeat across many different domains, not just drug discovery, where the old way of doing business can't keep up with the sheer volume that's being produced by AI outputs. And that's going to require, you know, you can throw AI in government and you can take a civil servant and give them AI tools. But to really capture the upside, we need to think about broader process reform. And that's where the executive action is limited. That requires Congress. to think, you know, what would be a new way of doing FDA drug approvals, let's say, and to think about your own area of jurisdiction or the topics you work on and think about how AI is going to directly impact that area—not just in a narrow economic sense, but what it means for the institutions that oversee that area.

**Jason Green-Lowe** | [19:52](#)

Okay. So speaking of upskilling in the federal government and getting them the resources that they need, what's missing either in terms of funding or in terms of training or talent recruitment to make some of these programs in the Executive Order a reality?

**Elise Phillips** | [20:17](#)

So, you know, I would endeavor to say... My other issue that I work on is competition. And what's worth noting within the competition section is that there's actually two uses of competition being used. There's competition from the global perspective and there's competition from the domestic perspective. And I think that's to append your earlier point, Bill, about making sure that we remain a global competitor. It's a helpful and distinct nuance that the Executive Order is making here. On top of that, there should be robust competition regulation on the domestic level. And what I consider missing from that is sort of a strong mandate for Congress to develop and pass bipartisan antitrust and competition legislation. On top of that, within that, there should

be robust resources provided to our antitrust enforcement agencies, which is woefully missing at this time. We are seeing large-scale developments when it comes to AI. And like somebody had also mentioned earlier, only the largest tech companies are capable of implementing this technology in a manner in which really creates a lot of the problems that we're concerned about and that we're seeing. And in order to address these concerns and in response to that, there needs to be an ask of Congress to provide these antitrust enforcement agencies with more resources. And I think that was something, in terms of the competition section that was missing, as well as some other things.

**Samuel Hammond** | [21:55](#)

Can I piggyback off that? You know there's ways that we could you know boost funding for different agencies and so forth But I think if you look at the FTC They have a health care division that's charged with policing competition over the entire health care industry so all the pharma all the hospitals everything And it's a 30, 29 person team of attorneys right so you got 30 people, 30 attorneys, that are in charge of policing the entire health US health care industry, which is like 18% of GDP. What does the day in the life of a FTC, a healthcare attorney, look like? It looks like subpoenaing 40,000 emails from a pharma CEO who's accused of misconduct and then manually combing through those emails to find evidence of that misconduct as part of discovery. With today's tools, those attorneys could put all those emails into Claude-2 or GPT-4, obviously something with confidentiality and privacy protections. But simply ask the model, you know, find me the five most egregious examples of misconduct. And, you know, there might be some prompt engineering that has to go into that. And even if it's not perfect, even if the attorneys then go, like, trust but verify, it will massively accelerate their productivity and grow their capacity to take on new cases too, to pursue the cases they're already doing more diligently. And I can guarantee you that the private law firms that are on the other side of that lawsuit will be using these tools. You know, Big Law is investing heavily in sort of the AI lawyer co-pilots and so forth. And so I think. This goes again to this idea that you have to fight AI fire with AI fire. There's a diffusion, there's an arms race going on between the companies, but there's also an arms race going on between the diffusion of AI into the private sector and people adopting tools more rapidly, and within government. And if our government is slow to adopt AI internally because of perceived risks or imperfections, then they're going to be basically playing with one arm tied to the other back.

**Jason Green-Lowe** | [23:56](#)

Daniel, what does the public think about increased use of AI tools within the government?

**Daniel Colson** | [24:04](#)

So it depends a lot based upon the particular application. So, you know, for things like writing news articles, I believe only 40% of the public thinks that that's a good use of AI versus 60% that is opposed to it. For things like driving cars, more like 75% of the public thinks that it's a good use of AI. In terms of usages within the government and sort of security critical applications—you know, things like usage of AI in warfare, AI making lethal decisions on the battlefield without a human in the loop—those are the areas where you see totally overwhelming

public opposition, 75% bipartisan opposition to allowing AI to make totally independent, lethal decisions.

Something else I wanted to comment on was just in terms of things that the Executive Order is missing, you know, In these areas of extreme high risk, I think there's probably certain more straightforward bans and restrictions that would be appropriate. One that I've seen discussion of is concern around open sourcing of powerful foundation models and the way that that allows proliferation of very powerful technology to America's adversaries. I think this is gonna be an important area of, I think, some amount of conflict, because I think particularly Meta and Mistral in France are very committed to a path of open sourcing very powerful foundation models. And I think there's a concern that that can easily lead to proliferation of very dangerous capabilities, most especially because when the weights are available, you can sort of arbitrarily modify the model to be able to remove any of the safety restrictions that are on the model. So something that you see with image models, which are a little bit ahead of text models, is that even though the model isn't trained on pornographic images, with \$500 and basically a home computer, you can modify the model such that it retains all of its very advanced image capabilities, but extends its capabilities to a domain that it wasn't originally designed for. And so something that you'll see is, you know, if you have a \$10 billion trained model that carefully excludes, you know, biological training data, you'll keep all of the advanced general-purpose capabilities. And with a couple thousand dollars, you train it on biological data, and now all of a sudden it has the dangerous capabilities you want to remove. I think that creates really difficult proliferation dynamics around dangerous capabilities. And so, you know, I think consideration of certain more straightforward restrictions is something that I think would be appropriate as Congress looks towards legislation. That was something that wasn't really a part of the Executive Order.

**Elise Phillips** | [27:33](#)

Just to offer a different perspective and push back on that just a little bit as an organization that advocates for Open Internet and acknowledges a lot of the realities of open source. From our perspective, open source in a lot of ways can be beneficial for competition because what we don't want to see is only these large technology companies with these capabilities. We want to be able to diversify the market and provide robust competition to be able to have an alternative. If a company issuing an AI model isn't doing what we would like them to do and what we would like them to perform. And there has been some contrary evidence as well that, you know, open source can be good for data security and prevent vulnerabilities from existing in the first place because we have access and we are aware (we're all aware of sort of the specific open source code that's being utilized). So that's just a different perspective. We, as an organization, tend to see open source in a lot of ways as an inevitability. And so we need to, therefore, have a more cautious and appropriate approach to how can we mitigate some of the harms when it comes to open source? How can we promote competition? How can we diversify the market in a way that there are new manifestations and new use cases that can provide healthy alternatives?

**Jason Green-Lowe** | [28:48](#)

And Elise, is there anything in the Executive Order that would help with that goal?



**Elise Phillips** | [28:53](#)

I wouldn't say so. And that's just another way in which I think the competition section could be improved. There could be some more nuances in terms of the dominance of global companies. There could be just acknowledgement of potential conflicts of interest, sort of the strengths of a lot of companies that are vertically integrated. and the power that they have over different supply chains and streams.

We've seen, I think as of now, a lot of companies that are implementing a lot of these AI models. It's just the big ones. And so they're sort of using that access to this technology as not just a way to further their existing dominance in their Big Tech competition space in their specific markets, but also create dominance in the AI market itself. And so in a lot of ways, we see open source as a way, or it can be a way, to diversify and provide opportunities to smaller entities who are looking to expand and scale up.

**Jason Green-Lowe** | [29:56](#)

Okay. I want to throw the panel open to questions from the audience. So please raise your hands. If you get called on, briefly introduce yourself. You can say if the question is for a particular panelist or for anyone who's interested.

**Audience Member 1** | [30:12](#)

Part of this is going to be presented as kind of a race between government and communities. But how does the consumer level, the individual level, one, kind of prepare for this kind of revolution? How would we point our constituents to being informed as this race is going on?

**Jason Green-Lowe** | [30:36](#)

So what is it that constituents need to know?

**Audience Member 1** | [30:39](#)

I think so. Yeah.

**Jason Green-Lowe** | [30:41](#)

Okay. Bill, we haven't heard much from you. You want to take that one?

**Bill Drexel** | [30:48](#)

Um, sure, though it wasn't one intuitively I would have. I think, honestly, it kind of goes back to what I mentioned in my first response, that the space remains so much in flux that it's hard to recommend anything that's going to be viable for the next two years. I think that maybe, I mean, I think the lowest hanging fruit that average people would want to know about are some of the misinfo, disinfo, and kind of generative AI things around phishing. So there's a cybersecurity element where we're likely to see a lot more kinds of manufactured texts that can be used to manipulate individuals and also especially that can be used to try to get their financial details or various other means. There's good reason to believe that the models that are already here can really greatly expand the scale and quality of a lot of attacks that might affect individuals across the nation. That's one place I would start, but I'm sure there are others if anyone wants to add.

**Samuel Hammond** | [32:08](#)

I think sort of public awareness is going to be a really important part of this. And obviously knowledge is unevenly distributed in society and we still today have vulnerable populations like the elderly that get targeted by scammers in India and other places. Those call center operations are often very labor intensive, and we're going to move into an era where that's no longer the case, where you can have basically a server rack that's running thousands of robocallers that have been fine-tuned to be ultimately persuasive. And even at the congressional level, I think there's going to be a moment in the next couple of years where you pick up a call from a constituent and you're not really sure if it's a real constituent or just a model that's affecting a Louisiana accent.

And that goes again to this arm-race dynamic where we're going to need tools in the legislative branch to be able to screen for and verify for human identity and so forth. At the consumer level, it's going to be, I think, very confusing because it's everything everywhere all at once. And so there's going to be a lot of great things happening with a lot of just very strange and weird things. Because consumers are sort of the first to adopt, they don't have any constraints, right? Like if maybe the law says that a doctor's the only one that's able to give you a diagnosis, but if you go to GPT-4 and put in your blood work, it will give you that diagnosis. So there's going to be all kinds of ways that these tools break down existing barriers and break distribution—sort of distribution monopolies, like health care distribution monopolies or content distribution—and people are just gonna go directly to the source.

And that's actually, I think, a good thing, right? Like, you know, maybe these systems will hallucinate now and then, and so maybe they wouldn't get FDA approval as a medical algorithm or something like that. But if they still provide genuine value and they do it at a very, very low cost, like often free, people are going to try that out and they need to be aware of its limitations so they can be sort of informed users. That's gonna come through trial and error. And I think there's also gonna be different dynamics that play out over time where, you know, during the pandemic, for instance, when everyone was being forced to do remote schooling, there was a 2 million student increase in homeschooling. Right, the pandemic as an event helped coordinate people to realize, oh, we can opt out of the system. Now what happens when we have, you know, AI tutors that you can load into your browser that are as good as any tutor and ultra-effective. People who were thinking about doing homeschooling and are thinking about exiting the system suddenly have this option that they can just choose to use and coordinate around. Those kinds of effects are going to roll over society in a kind of way that no one's really controlling. It's going to be self-organizing. And the people who are sort of doing that are going to have externalities on the people who choose not to do that.

**Jason Green-Lowe** | [35:19](#)

Thank you, Sam. I just want to take some more questions. It's a great answer.

**Daniel Colson** | [35:22](#)

Could I add one thing? I think the labor impacts will be unevenly distributed. And I think educating people in particularly at-risk areas, you know, if you think of today, you know, like transcription and translation services, it's just like... it's not really going to last. But that was a lot of people in the past who worked in those areas. So I think educating constituents that are at the most risk for labor disruption as well.

**Elise Phillips** | [35:53](#)

I'd like to add one point as well, just as a former staffer who worked on these issues, you know, it's helpful to materialize it in groups. How does this impact consumers? How does this impact small businesses? How does this impact workers? And for consumers, obviously, when it comes to dominant markets who are controlling every aspect of the supply chain, that's gonna impact prices. We saw that during the pandemic on Amazon and digital platforms. That had marketplaces where folks didn't know where to turn to purchase PPE and we saw prices skyrocket. AI has the potential to expedite and increase self-preferencing, so that makes that difficult; it makes it harder for consumers to not only find an alternative, but when they're using a platform that they have no choice to use because there's no market alternative, they have to deal with these prices. Same thing when it comes to small businesses and digital platforms. If they only have access to using a specific digital platform to sell their products, and then an algorithm is scraping the internet and finding prices, and mandating that they provide the lowest price available. That's going to create a dynamic in which they have to drive down their own prices for their own products in response to that, and therefore they're making less profit. And also for workers and laborers as well. It has a similar impact when it comes to worker surveillance or wages, et cetera. And so there's a lot of ways and a lot of anecdotes and a lot of organizations that can materialize these harms.

And what you can do is draw anecdotes from that and communicate that to those constituents. If you're confused why prices are skyrocketing, this is why, and this is what my member is doing about it. This is the legislation that my member is supporting to address these issues. And this is what you can do to mitigate those concerns and educate them in that way. I hope that's helpful.

**Jason Green-Lowe** | [37:43](#)

Great, okay, who else? In the back, yes.

**Audience Member 2** | [38:18](#)

I wonder what ratio would be geared towards different groups?

**Bill Drexel** | [38:30](#)

All right. One thing that I think is interesting about the Executive Order on this topic is a notable absence. It has something for everyone, except defense, more or less. So the Executive Order has more or less indicated that in as far as the government or this administration is looking at AI issues, they're looking at a broad range of domestic issues trying to balance the interest groups you're discussing. And then on the side, somewhat independently, they're looking at the defense applications and pursuing a kind of semi-independent trajectory on that based on what China's doing, what other kinds of allies and partners are doing. And also, I mean, partially because the

U.S. government, the military has actually been relatively ahead of other militaries, both on looking ahead for capabilities and also looking at safety and responsibility questions. I think personally, I think that's a good approach. I think it's worth having a domestic debate about domestic issues, but when it comes to defense, there are just some things we have to do and it follows a different logic and has different implementers. So I was pleased with that element of the Executive Order as well.

**Jason Green-Lowe** | [39:58](#)

Thank you, Bill. It's a fair question. It also has a tendency to devolve into stump speeches. So if other panelists want to take a crack at it, go ahead, but try to focus on the balance of things in the Executive Order and not why your favorite cause area is the most important.

**Samuel Hammond** | [40:14](#)

Well, I was going to say, I don't think there's actually a huge right line between the consumer interest and our national interest in this case. I think one of the fallacies in the debate with the competition with China is that we need to just hit the pedal to the metal and deploy this technology as quickly as possible because if we don't, then China will. When actually China's taking a very different approach. They're being extremely conservative about who's able to train these models, how they're deployed. They have a draft law that basically says any AI model that undermines the unity of the nation or socialist values is prohibited.

And the reason China's taking that approach is because they saw how the internet had these double-edged effects, right? If we were having internet safety discussions in the 2000s, we would have been talking about identity theft or consumer protection. You wouldn't have anticipated that the internet and social media and mobile would help lead to a legitimacy crisis across Western democracies and forge the Arab Spring and weaker states. And China was watching that, and they saw how the internet and open access information, which we value in the West, was destabilizing to the regime. And so they had to crack down and build a surveillance state around that. And they're taking a very similar approach on AI. And I think we need to incorporate some of that lesson ourselves to think. To what extent is going full steam ahead, meaning that we're going to dominate the technology and create a new American century, or if we don't do it sort of carelessly, lead to our own internal destabilization?

**Daniel Colson** | [41:52](#)

One further point on this is there's obviously a lot of uncertainty about how AI will develop over the next few years and a lot of disagreement amongst people who you might imagine would know. You know, I think the priorities today look very different if AGI is a 2026 technology versus a 2046 technology. That's a very difficult question. I think there's a lot of very not crazy people that say that AGI could be a 2026 technology and a lot of not crazy people that say the opposite. And so I think in that sense, the Executive Order was, I saw it as generally being mostly preparatory in the sense of trying to gather the information necessary to be able to deal with the technology and to set appropriate rules, but not doing very much to set those rules today. And I think in an important sense, I think that that tees up Congress for its own action.

**Elise Phillips** | [43:05](#)

Yeah, I mean, I would agree. I mentioned this earlier, but, you know, at least within the realm of competition, there's this nuanced balance between remaining globally competitive with other countries that have made strides in artificial intelligence while making sure that there aren't any net negative effects on consumers and small businesses domestically. And so you'll see, particularly within the competition section, a clause on immigration and how we're addressing the brain drain that tends to happen in the United States and making sure that we can retain talent. And so if you're speaking to people that have that specific interest when it comes to immigration, that's an angle that you can speak about when it comes to developing and prioritizing and making sure that we have the hardware and infrastructure. This Executive Order is appending to something that passed like the CHIPS and Science Act to make sure that we have sort of the processing capabilities to manage this. And that's a way that you can talk about AI in a way for somebody who's concerned about that and jobs as well.

And then from the consumer protection perspective, just the mandate of making sure that federal agencies are paying attention and scrutinizing companies that are implementing these technologies. So those are just different angles in which you can talk about the Executive Order and different nuances that you can address for folks with varied interests.

**Jason Green-Lowe** | [44:31](#)

Another question? Anyone working on possible oversight, possible legislation, possible hearing related to the Executive Order? Have a thought or a question about that? Oh, in the back. Okay, great.

**Audience Member 3** | [44:48](#)

Workforce preparation, building an AI understanding of future population employees.

**Samuel Hammond** | [45:09](#)

So there's stuff in there about workforce. There's, you know, DOL do a report on the potential labor impacts, so on and so forth. There's only so much you can do through executive action.

One of the sort of legacy issues that we have in the U.S. with our workforce system is it's, we have WIOA and it's a very broken system to say the least. Often implemented at the county level, often done through these workforce boards that have local union members, local community college, and maybe some business person, and then, everything is going to be some college course that doesn't really do anything.

And so I think the whole system needs to be reimaged from the bottom up. And I think AI can actually help retrain people and expedite upskilling. But it's not going to be done through sitting in a community college room. And I think we need to (folks working on labor in particular) think about a broader comprehensive reform that you simply can't do through executive action. The executive action includes a lot of silly things where it says like, you know, we're going to have the NSF invest in 500 new ML engineers to increase our machine learning talent. And it's like, I don't think the NSF is going to move the needle on our machine learning talent. I think the

private companies are doing that in spades and people just do it on their own volition. And so it's a lot of tinkering on the margin. And I think what we really need to think about is we're entering a new era, right? In the same way that the Industrial Revolution required many new institutions. Like totally new; unemployment insurance was invented, right, in 1955. We're gonna have to need new systems like that, and it's gonna require legislation.

**Elise Phillips** | [47:03](#)

Yeah, I mean, absolutely, I've said this and emphasized this a lot, but you know, there should be strong regulatory oversight for these capabilities and these technologies. And then something that I just felt important to note, because it was mentioned earlier, you know, concerns about corporate capture. You know, there should be, as we grow and we develop, when it comes to regulatory oversight, we should draw the balance between recruiting talent and knowledge from the private sector of folks that are knowledgeable about these technologies and therefore are effective at crafting legislation and having agency expertise to advise some of our political leaders, both in the legislative and executive branch, while also mitigating any potential anti-corruption concerns that come from like a revolving door dynamic. And so if your bosses are working on anything when it comes to regulatory capture or anything that comes to strengthening regulatory oversight for a lot of these things, I would encourage that.

I do know that my former boss has a regulatory capture bill. I'm gonna plug her on her behalf, shout out to her. It's called Stop Capture Capture. But there's a lot of different ways, legislative, that you can address sort of a myriad of issues that these technologies present. I kind of want to go into that. I mean, what can congressional leaders do to ensure that things like order and other priorities regarding AI are funded or can attract people from the private sector to come join us?

We have to make these desirable jobs. And there needs to be greater funding so we can hire talent and people stay and retain. A lot of people in the executive branch and in the legislative branch get burned out, as I'm sure you guys are aware. And so we need to provide more resources to these agencies. And we need to provide a healthy work environment which means great benefits, higher salaries, and incentives for people to come and stay. And then on top of that, there needs to be some ethical restrictions as well when it comes to lobbying, when it comes to anybody that may or may not be negotiating laws and or regulations that would be beneficial to a potential employer in the future. I hope that that's being understood as well, because that is a concern.

**Samuel Hammond** | [49:33](#)

AI as a general purpose tool, it's come first for writing your homework or for some of these knowledge jobs. I think bureaucracies in particular, corporate bureaucracies and public bureaucracies are incredibly exposed to this technology. OpenAI put out a research paper last year on GPTs as general purpose technologies where they estimated the labor market impact. And legal secretaries, auditors, these were occupations that were 100% exposed.

And to the extent that a lot of what the federal government does looks like basically being a kind of fleshy API, right, there's huge potential to increase the capacity of government without spending more money. And maybe there's some upfront cost to implement the technology, but even that, Microsoft's the biggest IT vendor in the US government, and everyone is running Microsoft Office on their computers, and they're all going to have GPT built into it. And so there's going to be a day where, you know, the person at the Bureau of Labor Statistics does a jobs report every month, is showing up to work and just loading a CSV file and saying, you know, write my blog post. Find the five most interesting trends. And so there's gonna be a need to reallocate labor within government, right?

There's, you know, 2.5 million in the federal workforce. You can do more with less with this technology. And it's gonna be a question of how it gets implemented. and whether these tools are able to be adopted or whether we need to actually bring in tons of specialists to do that implementation. And that can be done in a scalpel kind of way with special hiring authorities and so forth.

But going back to the oversight question, I think this is one of the biggest missing pieces of the Executive Order is how AI could be a massive boon to oversight, right? Because most of what the government does now is through machine readable text-based formats, emails and everything else. And Zoom calls are all transcribed. And so that's in principle something that can now just be audited, right? And so do we move, like when do we move to Inspector General GPT where we have all this stuff being able to, you know, being fed into a system that can generate reports to Congress instantly instead of, you know, in 180 days or whatever the norm is. And that could massively increase federal oversight, not just sort of in a trivial way, but like being able to kind of real-time monitoring of what the government is doing.

**Bill Drexel | [51:59](#)**

One further thing on that, maybe slightly adjacent to the point, but I think still relevant. This was referenced earlier, but I think there is really a big role for Congress to play in helping to expand our pool of AI talent, particularly from abroad. And anytime you mention the immigration word in Congress, immediately the red flags come up and you say, good luck, whatever. So I acknowledge that it is what it is. But at least with... I tend to think that we should. In terms of the numbers we're talking about, it's somewhat negligible to the broader immigration question. And to the extent that we can, we should separate those issues. I like the language of Sam's colleague, Jon Askonas at the FAI, who calls it weaponized brain drain. I think that that actually is the best path forward. We really need to give AI engineers from China and other countries, particularly adversaries, ways to stay in the United States, which they want to do. It helps us and it cripples them in an extremely important technology and it builds our pool of talent that we need in government and we need in the private sector and we need in defense. Maybe not putting those particular individuals in defense, but you get my point.

**Samuel Hammond | [53:31](#)**

Just to add to that, there's been some folks who've done sort of network analysis of the AI landscape in China, and there's probably like 100 engineers that if you gave them green cards

would decimate China's AI industry, or at least set them back substantially. And those people, maybe there's security risks, so let them do things that aren't sensitive if they're willing to come to this country and do the strategic brain drain, as Bill mentioned.

**Jason Green-Lowe** | [53:59](#)

Yeah, let's take one more question and then we'll open it up for informal off-the-record mingling.

**Audience Member 4** | [54:05](#)

I'm just curious if like AI or the Executive Order is being used for anything like space exploration or something like that. I just kind of read recently that, a normal human fighter pilot and certain jets can only withstand certain G-force situations and stuff like that. But once they start deploying AI into those kinds of environments, things will change. Are there implications for that, for like space propulsion systems and stuff like that?

**Samuel Hammond** | [54:36](#)

There's nothing particularly in the Executive Order, but it does have huge implications, right? So wouldn't it be great if when we start terraforming Mars, that we just send up robots that could build the base before we arrive? Because the first Mars mission is going to be probably a coin flip in terms of surviving, right? And so there's going to be like, humans will colonize space via robots. But the Executive Order doesn't really bear on that at all.

**Jason Green-Lowe** | [55:11](#)

All right, that was a fast one. One more question.

**Audience Member 5** | [55:27](#)

So the question is about using AI tools for government. The question of whether we're allowed to use basically this tool to create code. It's just being left up to the agency's office. How is this going to get sorted out? Like, who's going to fix this? Is this something that should be on the agency level? How should this work?

**Samuel Hammond** | [56:03](#)

I have thoughts on that. So I think one of the things that's still a big open question to me about whether it's a good idea or not is both this Executive Order and the follow-on one from OMB, they prescribe creating these new AI governance boards at each agency and appointing chief AI officers at each agency. And then there's a system to sort of bubble up to OMB. And I worry that that's going to be very duplicative and fragmenting. And every one of the discovery boards are doing their own independent risk analysis on tools like an Anthropic model that are general purpose and probably going to be used across government. It would be much better if those things were approved once and then any agency could know that this is a model that's safe to deploy.

The way we've done that with cloud has been through the FedRAMP system where it was addressing a very similar problem where in the early days of cloud all the different agencies had very different procurement policies and very fragmented and they created FedRAMP so they



could just have one portal you go to and say these are all the cloud services that are approved, that pass the NIST cybersecurity rules, and so on and so forth. We need something like that building on top of that central system for AI, where there's going to be bespoke tools that agencies need for their specific purpose, that they're going to procure on their own and need specific oversight. But for these general-purpose tools, there should just be one window and have that presumptively approved. And then within agencies, there can be more fine-grained, these are the uses, you can use it, and these are the limitations and so on and so forth. But it would be a drastic mistake if we just backtrack on the streamlining that we've already done.

**Jason Green-Lowe** | [57:52](#)

All right, well, thank you very much to our panelists for coming out and educating everyone on this incredibly important topic. I will be here for a while if anyone wants to chat and I hope some of them will be as well.