2023 TPI Aspen Forum
What are the Major AI Policy Questions

Scott Wallsten:

Okay. We'll get started with this next panel, which are What Are the Major AI Policy Questions? So just to set it up, do I control the slides here? I do. Okay, great. So I just want to give a very, very quick setup. Okay. Everyone's talking about AI. You can't go anywhere without hearing people talk about it, including here. So we'll just set it up a little bit. We know that generative AI's growth rate is faster than just about anything we've seen before. This slide shows the days to 1 million and days to 100 million users based on sources of ... I don't really know exactly how accurate it is, but you know that the general point is true. Then yeah, good. Also, it's getting better, at least better at taking tests with GPT-3.5, GPT-4. We've seen all the stories.

So that's made people excited and concerned, and you can see the way policymakers have reacted. That big bar is the number of state bills that are pending. A much smaller number, of course, have been passed, but every state is thinking about laws in addition to the federal proposals and thoughts that we're going to talk about, too.

So I just want to set it up with that, and we have a great panel to talk about all of this. I'm just going to very quickly introduce them. We have Halie Craig, who is a senior advisor at the US Senate Committee on Commerce, Science, and Transportation. Danielle Li is an associate professor at the Sloan School of Management at MIT. Katie Shay is Associate General Counsel and Director of Business and Human Rights at Cisco. Hal Varian is Chief Economist at Google and Emeritus Professor at UC Berkeley.

So I want to start at the very highest levels and work our way down to things that might actually happen. So the first question I guess is what is the right framework for even thinking about AI, regulating AI? Do we want to adopt a precautionary principle, which is Europe's general approach to everything, or do we want it to be more along the lines of thinking of actual harms that you have to think, "Why is the harm? Is there something we can do about it that would actually make it better?" At the highest level, I think that is the question of how we're going. So who would want to start with that? Sure. Go ahead.

Halie Craig:

Hey, everyone. My name's Halie. It's good to see you all today. So from where I sit in Congress, there have been a flurry, I think, of legislative proposals, some frameworks, some actual bills to regulate AI. I think part of that is frankly because it's trendy right now, and with the growth of ChatGPT, people are really paying attention to AI. But AI is not new, and I think policymakers need to take a step back and realize that while this technology has made significant leaps in recent months, the potential harms stemming from AI are not something that just grew overnight.

When we talk about permissionless innovation versus precautionary principle, I think we already have the data on this. You can compare US internet companies to European internet companies. Europe took a precautionary approach in the early 2000s, and the US allowed internet to basically grow unfettered and largely free from regulation. We have the most successful internet companies in the world. In AI, one of our comparative advantages is that we still don't have that precautionary regulatory framework.

So over half of all AI investment globally is currently in the United States. It's a global imperative for the US, in my view, to come out ahead in the AI race. It's going to happen regardless if we

regulate or not. China's going to continue to innovate and try and grow their AI sector, and so for us to maintain our leading edge, I think we need to be very, very cautious about jumping to regulate or jumping to legislate.

Scott Wallsten:

So you don't think there's something inherently different about AI that should make us think about it differently than other technological innovations?

Halie Craig:

What I think is important is to be able to drill down, and we talked about this a little last night, but to drill down on what the actual risks are that we're talking about, because they seem very nebulous to me from where I sit, like, "Oh, misinformation or bias." Well, those things already exist. They exist regardless of whether or not AI generates them, whether a person is biased, whether it's a classifier, an algorithm that is misfiring or is maybe not calibrated correctly. So putting AI in a unique category and regulating it as such I think misses the mark in terms of actually trying to get at these distinct harms or issues if those are things you are concerned about.

Hal Varian:

I could add to that sentiment a little bit, but first I have to warn you that not everything I will say will be true. I call that the Hal-lucination as much, and I'm trying to recruit Halie for the same principle here.

So this is a place where I think you really have to think about the mantra in the Silicon Valley. We overestimate what can be done in a year, and we underestimate what can be done in 10 years, because if we look at how rapidly things move in the media world, this is not a danger that's luring over us in the current time period of a year, two years, three years. This is something you're going to look at in the long run, and I think if you are patient, you'll be able to identify the true harms or the true places where regulation is going to make sense. I think it's really too early to determine that now.

Katie Shay:

Thank you, Scott, and thank you to the organizers of TPI for putting this event together. It's so wonderful to be here. So as Scott mentioned in the introduction, I lead the human rights function at Cisco. Halie, your comments about how these potential harms are not new really resonates with me. I think they just show up in slightly different ways as technology continues to evolve.

So how we're approaching this at Cisco is to bring a human rights approach to this and look at across our technology portfolio, how could our technology impact people, whether it's technology that we're using internally for business operations or features that we're putting into our products themselves? So from a human rights perspective, I think a lot about the right to privacy, and we have a lot of regulation that we're already subject to that governs how we protect data, especially sensitive data, and how we think about that from a transparency perspective as well. Same from a security perspective, and bias is a new issue that we haven't confronted with other types of technology, but certainly is really top of mind when it comes to AI.

So we've been doing human rights impact assessments at Cisco, and we've now adapted those to a new type of impact assessment we call a responsible AI impact assessment, where we look at potential risks that could occur in the technology portfolio from that perspective, and then our team presents mitigations to our engineers who can then go forth and implement those.

One of our key focuses within that also is how are we communicating to people who could potentially be impacted about the fact that they're interacting with AI and what they need to know in order to make good decisions for themselves about how they want to interact with that? So I think from a regulatory perspective, going back to Halie's comments, I think we need to give space to this technology to grow and adapt and develop and really leave the regulatory focus to those highest-risk use cases.

Danielle Li:

Hi. Thank you so much, Scott, for having me. So I think rather than think about regulation as what we can do, what we can't do, I actually would disagree a little bit and say that I do think that AI works in, I think, a meaningfully different way than other forms of computerization or information technology. So if you think about when computers entered the economy in the '60s and the '70s and the '80s, the way that traditional computers work is that you have to give them instructions, and you tell them what to do. You tell them how to add these numbers together. You tell a robot how to put this piece with that piece, and it does it. So it became good at tasks for which we could say, "Here are the instructions."

AI works differently in the sense that you don't provide instructions for how to use ChatGPT. What it does is it learns from examples. So you take a lot of training data, and training data can be defined in lots of different ways. Based on those examples, it learns, and new generative AI creates these behaviors. So this is not as common about regulation, but more about data. I think it's worth it for policymakers, managers, people to think about data infrastructure and how we affirmatively try to create data that provide more useful examples for us to learn from.

So let me think about two examples. So one is I've done some research in hiring. So increasingly, algorithms are being used to screen job applicants. There's lots of companies that use this information, a lot of companies that provide algorithms to sell to other companies so that they can use it. Usually what you do is you just find the data that you have, and that basically means that you'll get data on someone's resume and you'll get data on whether someone received an interview. Usually, companies keep track of who they interview and who they hire. But they won't really link the data on how well people perform on the job with the applicant data because that just requires them to have data systems that are better.

So you end up with a lot of algorithms that are being trained on the data that's available. The data that's available is whether you're interviewed, and so you end up basically training on people's existing behavior. So this conversation about bias, it's not just algorithm. That conversation actually should be a conversation about the bias comes from the kind of data we use. If we made it easier to link outcomes data, like actual outcomes that we care about, like how well someone does on the job, with the input data, we can actually build better algorithms.

You see this in healthcare settings, where it's really hard. You can imagine there's lots of data that we can use to think about what medical treatments or what drug treatments are going to be more effective, but it's very hard to link mortality data or healthcare data with characteristics data. There's lots of privacy issues and concerns with that. But I think that's an area in which rather than just regulating, we can think about how do we build infrastructure that enables more useful use of these technologies?

Scott Wallsten:

I think we should also come back to that as we discuss specific areas, and then also, of course, you can talk about your work on employment and connect it to what you're saying here. But before we dig more into the very specific things, just one more quick discussion about existential risks because that's always part of the discussion, not whether we think they are or not, whether

we think that the Terminator series is a good model for public policy, but we can think of these as low-probability, high-consequence events. These exist in lots of other things we do, too, in policy, not just in policy, in the world.

Again, is there something different about AI that should make us think that thinking of it as low probability, high consequence is not the right approach, that it's not such low probability, or should we think about it the same way we think about other issues in that category, developing nuclear weapons or energy or anything like that?

Hal Varian:

Well, we of course have gone through the same set of issues involving various kinds of weapons of mass destruction. I don't think there's really a lot of difference between who pulls the trigger. Is it an AI, or is it a person? This kind of terrorism, in my view, will exploit technologies of all sorts, as we've seen with the drone wars going on in Russia. So my view is there really is no difference. It's the same set of issues that we face with other powerful weapons and powerful technologies.

Katie Shay:

Yeah, I'm not well-placed to weigh in on whether there's an existential risk to this technology. But a lot of times, the work that my team's doing when we're evaluating human rights risk, whether it's in an AI product or any other product that Cisco makes available, oftentimes, the types of risks that we're flagging are low probability, high risk, where there's a potentially widespread impact even if it's low probability.

So what we do is we'll flag that and give recommended mitigations to teams so that they can try to reduce that even low probability or put something into the design that prevents that harm from occurring. This is really integral to the human rights approach. If you look at the United Nations' Guiding Principles on Business and Human Rights, which is what we constantly look to, it's a protect, respect, remedy framework. So you're conducting due diligence on a business activity, on a technology. You're identifying potential harms and then putting in place measures to prevent, mitigate, and hopefully not, but potentially remediate those harms if they do materialize. So I think that framework really can apply as new technologies come up when there is a low probability of harm, but potentially widespread impact.

Scott Wallsten:

Yeah, Halie. Go ahead, please.

Halie Craig:

So I'm not an engineer, so I feel like I can't talk too much about the emergent properties debate. I'll leave that to the experts. But I think we all saw the letter calling for the six-month pause on AI development that had, in our view, fairly alarmist language in it. Pausing AI development, just the United States, is probably one of the worst things we can do if we're worried about existential risks of AI. Even if these existential scenarios are very low probability, adversaries are still going to pursue them. Also, from where I sit, I don't think Congress is going to be the one that solves this problem.

If AI develops emergent properties or an adversary develops AI that can be deployed in a very adversarial manner, whether it's a new weapon or a bio-weapon or something like that, Congressional action is not going to be the thing that preempts that and stops it. That's just not how this is going to work.

But on the flip side, and we haven't gotten into this yet, but we also, I think, need to be focusing more on how AI can be deployed defensively and for good. So AI could be used to stop cyber attacks. Then outside of the defense realm, I've been thinking a lot about where AI can be very democratizing and very positive in terms of its influence on society, so advances in medicine, but also I used to be an education staffer. Democratizing access to quality education that is more personalized for students, that is really, really fascinating, and if we're not developing AI proactively in those areas, we're really missing out on the potential for growth in society.

Scott Wallsten:

Okay. As we go forward now, I want to talk about economic effects, ethics, and intellectual property. Also, I want to point out that we asked the four leading AIs what they thought the biggest AI issues were. It's a chart. You can look out there and see what it told us, which we thought were interesting in what they agreed on and what they didn't agree on. None of them said, "Nothing to worry about. No worries here."

So let's start with economic impacts. Maybe Danielle, you could start there because you've got work on productivity AI. So tell us a little bit about that, what you found, and then maybe connect it back to your comments before about regulation and what we might want to do or not want to do.

Danielle Li:

I think I'll just preface this by saying that there's very little evidence right now on the economic productivity impacts of both AI and generative AI specifically. So I'll talk about some work I've done, but I also think that that work raises a lot more questions than it closes.

So what my co-authors and I did was we studied the impact of the use of a generative AI tool in the workplace and looked at how it impacts the productivity and some measures of workplace experience. The setting was looking at essentially call centers, which are now called contact centers. It's basically actually one of the most common use areas of AI now in the sense that it's a place where the technological feasibility exists. So there's a ton of labeled data. Every time you have a conversation with a customer agent, there's a sense of, "What did you say? How did it go for you? Did you call back? Were you unhappy with it?," combined with large language models, lots of technological progress in that area, combined with a business need in the sense that this is an area in which workers turn over all the time. There's just lots of costs. People aren't really happy with sort of these kinds of experiences.

So we look at this, and essentially, I think a lot of the existing literature on technology in these workplaces has this flavor of if you're good at your job, it helps make you better at your job, and if you're bad at your job, it leaves you behind, so this relationship between technology and rising inequality. Our medium-term evidence, so six months-ish evidence, is actually not that. What we find essentially is that the way that generative AI works is that you get lots of examples on how people do their jobs, and machine learning technologies are designed to make distinctions like, "This is correlated with a good outcome. This is correlated with a bad outcome. I can make these predictions if I have this data."

When you take ML machinery and you apply it to data that's generated by human beings doing their job, what you are actually getting is you're getting an understanding of what kinds of behaviors people who are good at their jobs do versus people who are bad at their jobs. It provides a model of what makes people good at their jobs, including stuff that they do that they can't really tell you that they do. So there's this tacit knowledge.

So the result in the paper essentially is that productivity goes up a lot. In particular, almost all of the productivity gains are being driven by people who are newer and people who were at the

start less good at their job. So what actually happened is that if you're not that great at your job ... Halie, she's great at her job. She's not sitting over my shoulder, telling me how to answer a question, and I don't have access to that. But what the ML does is that it creates this model of her that now gets to sit next to me and help me do my job.

So if anything, we see a compression in productivity. So it brings newer workers up, and it makes the experience of work better. So we oftentimes think about how AI can just ... how it's going to destroy, how it's going to make the job worse. But actually, in a lot of settings, it's really unpleasant to be bad at your job. In call center settings, you get yelled at a lot. So we did sentiment analysis, and it seemed like people got yelled at less. They stayed on the job more. So your job was less unpleasant, so you are more likely to stay on the job. So I think this illustrates in concert with the conversation about negative impacts of AI that there are potentially large positive impacts.

What I would say is missing here in this paper, and it's a time thing, is thinking about how workplaces adjust to this because you can imagine a situation in which now people who are novice workers, they're given more responsibility. So the positive story, I think, is that someone with access to generative AI can do things that they weren't able to do before. Someone who wasn't able to code can now do some coding, and that creates job ladders that lead to economic mobility. So in some sense, this allows lower-skill workers to access higher-skill tasks, upskill themselves, and move up.

The other story you can imagine is that this in fact makes it easy for you to do your jobs without your interns and without your research assistants, so like the email story. We used to have assistants do our correspondence. But now email makes it so that we do our correspondence, and we don't have as many assistants doing this. In that world, you actually do end up with polarization in the sense that you have high-skill jobs. People doing those jobs, they don't need to have a lot of people working under them, but because there aren't people working under them, there isn't as much mobility, moving upward.

I think that that's the central question, the impact of how AI is going to change how responsibilities and learning opportunities get bundled within the workplaces. That's missing from this paper, and I think that that's something that we really need to understand to think about what the economic opportunities and harms are going to be.

Scott Wallsten:

This is great objective positive analysis, but if you're talking to somebody in Congress, like somebody next to you, for example, and thinking, "What does the paper mean for policy?," what should policy or should policy not worry about in terms of labor issues? I know that's a stretch from the paper, but-

Danielle Li:

Yeah. I think one benefit of being an academic is that you don't actually have to make the hard decisions on the spot. You can sit there and make sure all the evidence is right. I do think it is important to make it possible to run these experiments that Matt was talking about in order for us to make policy with better data. One of the things is that in this setting, I really do think we want to study, "What are these kind of bundling aspects?" before doing that policy.

I think even without that knowledge, we can play down the game tree in the sense that we know that some people are going to be winners and some people are going to be losers in a massive way. We don't necessarily know who that is in this moment. But we can think about once we figure that out, what are we going to do with people who have been left behind as a result of these technologies?

The other thing that I would probably feel comfortable saying is that people's fluency with AI, their ability to work with these tools to have it augment them, rather than replace them is going to be really important for your labor market success, and I wouldn't want to live in a world where access to that education is very inequitable. There's already a lot of evidence that if you don't get a good education, you end up falling behind. I feel like that these tools can amplify these things, and building this AI literacy into how we train people I think is important.

Scott Wallsten:

Hal, you've done a lot of work on [inaudible], and then also, Katie, you've all put so much into an ethical framework that includes a lot of things Danielle was talking about, so maybe you could talk about that. There are a whole bunch of folks that we talked about that relate to that. So yeah, Hal, go ahead.

Hal Varian:

Yeah, I just wanted to talk for a minute about the productivity paradox because we still don't see productivity increasing in the data, and the question is, "Why not?" My answer is that there's been huge productivity increase in technologies or in activities that weren't in GDP before the internet and are not in GDP now.

I have a few favorite examples. One is photography. So it used to be photos cost about 50 cents apiece for developing a film, and there were only 80 billion photos taken around the world. Well, 2015, just 15 years later, there were 1.6 trillion photos taken, and it's still increasing. Most photos are shared, not sold, so they don't show up in GDP. You have this huge change, which has had all sorts of ramifications in terms of people being aware of what's going on in other countries, in terms of police behavior, in terms of weather, in terms of all of these disasters that we can watch on TV.

There is a strange productivity paradox because once you have a camera, once you have photos taken, then your demand for maps, music players, flashlights, alarm clocks, all these things disappear. What happens is they're no longer being sold. They're pulled out of the consumer's basket. Why? Because they have this substitute, namely the mobile phone, that allows you to do all these things without any additional expenditure. That device, mundane as it may seem to this group, has been and will continue to have a huge effect on society and behavior. The earlier talk about people's attitudes and people's self-image and all this kind of stuff, the digital cameras are playing a role in that kind of behavior. You wouldn't have Instagram, you wouldn't have YouTube if it weren't for the availability of those devices.

One other example of things that were used before the internet and continue to be used but don't show up in the economic statistics, think about maps. Some of you in the audience are old enough to remember you had to fold it out and plan your journey, and all sorts of things went into planning to get from A to B. Those are now gone.

Scott Wallsten:

Spend the rest of your journey trying to refold it.

Hal Varian:

Well, that was my punchline.

Scott Wallsten:

Oh, sorry. Okay.

Katie Shay:

Yeah. So Danielle, your research is really interesting in the productivity space, and I think that's certainly something that companies in particular should be thinking about, like, "How do we upskill our workers and prepare our workforce to interact with this technology in a way that is responsible and augments their current skillset?" As the skillsets that we need change over time, that, I think, is a good area for investment, whether it's by companies or governments, to upskill the workforce as well.

This is something Cisco has been focused on for a long time. We have a program called Networking Academy, which is a global program where we train people in networking skills to bring them into this new technology and the new skills that people are going to need in order to participate in the workforce.

But I wanted to pivot a little bit to talk about how other technologies show up in the workspace that may not impact productivity, per se, but participate or bring about a better workplace experience. This is technology that Cisco develops that I'm sure you've all interacted with, whether it was Cisco or one of our competitor companies using AI. Virtual backgrounds is one of my favorites that uses AI as a privacy-enhancing feature. But my personal favorite is background noise reduction, which uses AI to augment vocal sounds and suppress everything else.

It's my favorite because during the pandemic, none of my colleagues could hear my baby crying as I was bouncing her on my hip or the construction project across the street or my neighbor's leaf blower when I was working from home. So it really made my work experience better to have that technology in place, and it relied on AI in order for it to work.

The reason I think it's important to talk about in this context where we're talking about regulation is that those technologies also rely on biometric data in order to be trained. We're seeing a lot of regulation focus on AI systems that use biometrics as particularly high-risk use cases, but using biometric information to train a system that produces a virtual background or makes my voice come across more clearly in a meeting should not be treated as high-risk compared with other systems that use biometric information to identify people who are potentially involved in criminal activity, for example. So I think it's really important in the regulatory space that we're very focused on the use case and not only on the data that's being used or exactly how the models are being trained.

Another thing that I think is really interesting to pull out from the virtual background example is around bias. So one thing that we learned internally when we were training the virtual background system was that we actually needed to put skewed data into the training models in order to get an equitable result. What our engineers saw as they were testing this in the early phases of the product was that women have more variation in their hairstyles, and so they actually needed to put more training data in corresponding to women's faces, hair, bodies in order for our hair to not be cut out. There's an outline that gets close to the head if it's trained even on 50/50 male-female data. So we needed to put in actually what most people would consider to be biased data in favor of females in order to get an equitable result.

So I think that's just important to call out as well as we're thinking about regulation around bias, that it may not be the obvious thing that we all think, "Oh, let's just make sure that the training data is representative of society overall." But actually, maybe we need something else in order to have that equitable result.

Scott Wallsten:

That's really interesting. That's actually related to Catherine Tucker's work from some years ago where to have an unbiased Twitter campaign, you had to spend more money targeting women

than men because women were worth more in the auctions. It's really interesting. I'm sorry. I interrupted somebody.

A quick question here we'll take as a separate question by itself because I thought I like it on Slido, "Are search engines and answer engines like Chat GPT competitors or complements?"

Hal Varian:

I would say they should be considered complements for the most part because the question is, "What's the goal of your use of the search engine?" Is it to resolve a fact, in which case a classical search engine probably makes the most sense, or is it to offer an essay or an overview on a topic? Then you might be using ChatGPT kinds of technology for that. So it's a question where what is your goal? How do you achieve that goal? How do these tools interact? Those are issues we're exploring now. Those are the critical issues for the coming few years.

Scott Wallsten:

Was it a good idea or a bad idea to mix generative AI with search engines when Bing made it as part of it, and now Bard gives you answers when you ask?

Hal Varian:

Well, as of now, I think that I view them as complements. That is, it wasn't a bad idea. It was a good idea, and we'll see how that works out, subject to change.

Halie Craig:

I think just touching on that point, one of the distinctions, and this is just coming from my own experience, playing around with ChatGPT and the AI-augmented Bing, but because the generative AI tools like ChatGPT do hallucinate and do create fake research papers, for example, I think there still is something to be said for a traditional search engine being able to check the primary sources of documents and that sort of thing.

So I do see them as complementary, but they are just different use cases. I'll use ChatGPT to create an initial draft of a letter and then edit off of that, but I'm still going through the primary source material. So from a research perspective, I think of them quite differently. Of course, that is all subject to improvement over time.

Scott Wallsten:

The papers ChatCPT thinks I've written look amazing.

Halie Craig:

Your Nobel Prize is coming, too.

Scott Wallsten:

Exactly. I want to move on to IP, but Halie, I meant to ask you, just to finish up the jobs question ... Sorry for the interruption of a little bit of a cocktail talk there. But Congress does think about these questions of how you might compensate the losers. We've done that for trade and other areas. So what are people talking about in Congress? What are the proposals? What are your thoughts about it?

Halie Craig:

So my supply sider is going to start showing. I actually view the situation we're in now with AI, and this has happened in the past, too, with every new wave of tech innovation, like the printing press or audio recordings for music, that sort of thing, but I actually do view this as fairly analogous to trade liberalization because fundamentally, it affects the division of labor within the economy due to increased comparative advantages. That will be domestically and internationally when it comes to AI and new advances in technology. With trade, obviously, it's competition across geopolitical borders. But the result is not pure job loss, and it never has been. But the effect is a change in job composition. This, I think, goes to your point about the research assistants and moving into email correspondence.

But there have always been proposals in Congress for how to compensate the losers from the shift in job composition, and those programs I think have had mixed benefits. The big one, of course, is trade adjustment assistance. There are very few people on that program right now, and it hasn't had great outcomes. So I think maybe more to Katie's point, the question is, "How can potentially the private sector get more involved in upskilling?"

When I worked on education policy, a big thing we saw in more rural areas of Pennsylvania, which is where I'm from, was community colleges working with local businesses to train people to work locally. We're going to see this play out very soon with chips and all the new fabs that we're constructing that we don't have construction workforces for or enough people to even operate those fabs. So that's going to be an interesting test. But right now, we have a labor market shortage, and so I think where AI can be helpful is probably more in the services space right now. But we'll see how it develops.

One stat I like from the trade space is when we originally signed NAFTA, Pennsylvania lost about 16,000 garment industry jobs to Mexico is the estimate, but our exports total to Mexico quintupled. So the benefits for other areas of the economy and for economic growth as a whole are really huge, and we can't miss that part of the picture. On the compensation side, the only other thing I've seen is an automation tax proposal. I think that was a Sanders proposal. Needless to say, it's a conservative office. We are not supportive of that. But I'm curious actually for others' views on it maybe after the panel.

Scott Wallsten:

Okay, and then we want to talk just quickly about IP issues. Then we'll take questions, and of course you can continue asking questions on Slido if you want.

So intellectual property, copyright, there's two sides of it, the one question, this debate, should training models be able to use copyrighted materials without attribution or payment? Then on the other side, who gets to own the output from generative AIs, or who's liable for that?

So on the input side, is an AI different in its learning from a regular person in that everything we know is based on all kinds of things, copyrighted or not copyrighted, and we don't walk around, citing all the things that made us who we are? So even the case where Sarah Silverman is a plaintiff against OpenAI for this, not to criticize Sarah Silverman, because I think she's hilarious, but she presumably has built her work on Lenny Bruce and all kinds of other things and doesn't cite it constantly. Is that a wrong way of thinking about this? Is there more to this question of using copyrighted materials as inputs? What are other ways of thinking about it? What do you all think?

Hal Varian:

Yeah, I have a feeling here of deja vu. It's YouTube all over again.

Scott Wallsten:

Oh, I thought you meant it was because we talked about it last night at dinner.

Hal Varian:

That, too. That, too. I think I appeal to the coast theorem here. The coast theorem here says if it is their profit to be made, people will figure out some set of contracts or arrangements or practices that enable that to happen. That's indeed what happened in the YouTube case. They worked on this for several years and ended up with a compensation scheme that was a reasonable compromise to the several different views that were out there. So my view is you will see a normalization. It would take a couple years at least. It took several years in case of YouTube. But we have a model then that can be applied in that situation. So I think we just will have to put up with this negotiation for a period of time, and eventually, it will be resolved.

Danielle Li:

So in the paper that I talked about with all that productivity increase, like the low-skill people become much more productive, and the high-skill people don't become productive, the reason that happens is that the AI is being trained on the behaviors of the high-skill people. So if I'm a worker, I'm generating an example of how to deal with this particular angry customer or to deal with this particular technical problem. That is in a traditional world my productivity. But in an AI world, how well I deal with this particular customer isn't the main aspect of my productivity anymore. It's how that example gets fed into the AI system, which then has everyone else at my firm talking like me. So there's a way in which I am actually very much not compensated for the productivity impacts of my contribution.

In settings where it's literally Sarah Silverman or specific people, like, "I made this film," they're larger entities. In a lot of firms, what's happening is that there's just lots of just individual workers, and if you take one of them out, you can have someone else who's also good at their job. But there's a way in which I think AI systems do change the notion of what productivity looks like and what marginal impact looks like. I do think you need to provide incentives for people to generate those examples that keep these systems learning.

Scott Wallsten:

Halie, go ahead.

Halie Craig:

I was just going to add that I think there's a huge policy tension here, and I don't have the solution. But it is really interesting to think about because one of the biggest incentives that companies, including foreign companies, have to come and invest in the United States and for businesses to develop products in the United States is our IP, our very strong IP regime.

There's a reason that a lot of biologic drugs are developed in the United States, and I know that's a controversial topic on the IP side. But India is not doing that, for example. It's the United States. So we need to be able to maintain the strength of our IP protections, but also, if we want to develop stronger AI and AI that's more nimble and more accurate, that AI needs access to high-quality training data. A lot of times, that will be something that is potentially copyrighted or has some IP implications, and so it's a lot to think about and something I'm also interested in hearing from other folks on.

Scott Wallsten:

Yeah, go ahead.

Hal Varian:

There are sectors of the economy where you do not see intellectual property protection. It's customer norms in the group that facilitate the innovation. So for examples are in the US, fashion. Not true in Europe, by the way, but fashion cannot generally be patented or copyrighted. Recipes for food, cooks' trade recipes, and the point there is to acknowledge that this recipe came from so-and-so, but there's no particular intellectual property behavior, and jokes. You take Sarah Silverman, you can't really assert, "I have a copyright to this particular joke," unless you have exactly the right wording and it's using exactly the right context. Other sorts of entertainment, same thing. After all, we all learned in school about copying scientific theorems or geometric proofs or any of these other things. So it can be controlled or can be managed or can be made efficient both by custom and by law.

Scott Wallsten:

Questions? Yeah. Here comes a microphone. They're both running. We'll see who gets to you first.

Audience Member 1:

Thanks, Scott. Halie, I completely agree with your assessment that the existing state and federal sectoral regulation cover the risks of bias and harm, but we should always look for gaps. The Biden administration in July 25th came up with what it thought was filling the gaps with its voluntary commitments that generative AI would be audited to make sure that the answers coming out of it wouldn't generate societal harms. From the standpoint of the administration, what they think is disinformation or societal harms might be different than what others think. But I'd be interested in the panel's reaction to Biden's new voluntary commitments.

Halie Craig:

Okay. I can start. Yeah, so we re-reviewed these, and I think what stood out to me, at least on the first page, is the line, "These voluntary commitments are only a first step in developing and enforcing binding obligations." Since these voluntary commitments were agreed to, we've already seen folks on the outside end in the administration calling for these to be binding or obligatory or otherwise codified, and that very much terrifies me. We saw the same thing with the NIST RMF framework during NDA. Our committee was asked to clear making that mandatory, and we were able to stop that from happening.

It was just shocking to us because, of course, it was designed to be flexible and voluntary as a framework, but this is how regulatory capture works, to be honest. It wasn't really lost on us that the companies that signed that voluntary commitment pledge were the largest players in this space, which is fine.They're trying to protect their own business interest, but we don't see startups signing onto this. It is so early on in terms of the growth that I think we're going to see ... I'm an optimist here, but the growth that we're going to see in AI to be committing the entire industry to a set of standards that only the largest players potentially are able to meet.

Once you put that regulatory burden on almost every player in the industry, even if you have carve-outs for lower-risk use cases or smaller players, they're going to grow. What becomes risky is that goalpost is going to move. But once you put that regulation in place or these commitments in place, that's a cost that subtracts from the potential productivity and the potential for innovation. I don't think we're at a place where we're ready to assume that cost.

Scott Wallsten:

So you thought what they put out initially was the right track, but then quickly went off the rails? Is that-

Halie Craig:

No, I didn't like what they put out.

Scott Wallsten:

So you didn't like it right from the start. Okay. I misinterpreted, and I was surprised.

Halie Craig:

You know this, Scott, but I worked in trust and safety at Meta for two years on content and algorithm policy. So I did see the trust and safety debates happening inside the company when it came to speech on social media, which is why the parts of the voluntary commitments and the AI Bill of Rights, which is really about all algorithms, really terrifies me, the parts that touch on speech. Generative AI, of course, is creating speech and content. To the extent that there is a push either on the Left or the Right to look at outputs or look at inputs and try and prescribe regulations on top of that, that has real speech implications. So I am very wary. We can talk about cybersecurity issues or defense, but I'm very concerned about a potential regime where a federal regulator has the ability to dictate what an AI can and can't say.

Scott Wallsten:

Katie, I don't want to put you on the spot, but I want to put you on the spot because the ethical approaches to AI that Cisco put out, you talked about some of this. So how did you all feel about it?

Katie Shay:

Yeah, I think we're doing a lot of what regulators are talking about internally. But the difference is that we can apply it in a way that really makes sense for our business and our culture and our structure and all of that. We have a responsible AI framework that has principles and controls built into our design methodology. We have an internal oversight committee that has representation across business functions, with heavy representation from engineering and legal and others.

I think, Halie, the concerns that you're raising really resonate because we just see tremendous opportunity here with AI. I think it's really important that there's space for innovation to continue and especially for low-risk use cases to have that space to develop. We're invested within our company in doing that in a way that's ethical and responsible, consistent with our brand, and in a way that preserves consumer trust. So I think all of that can be done without really heavy burden regulation on companies.

Hal Varian:

But the issue is, is it appropriate to be censoring speech in these dimensions? There may be cases where I think it's important that people have the right or the ability to present opinions. I agree with Halie that it's premature to consider these issues. Now, they may arise with children or certain kinds of questions and things that are clearly harmful in the future, but we won't really know that until we run into these problems. So I have to say I think you want to proceed pretty cautiously down that road.

Scott Wallsten:

One more question. Wait for a microphone, whomever's closest. Sorry. Otherwise, people online can't hear.

Audience Member 2:

So I'm a big fan of innovation and agree with Halie that we also have to look at the AI for good, and the other one I would add on your list is climate change, using AI to try to address that crisis. But I want to just challenge the premise that the lack of federal legislation today equates to opening up the room for innovation because if you look at the landscape today, as you alluded to at the beginning, Europe is more likely than not going to pass sweeping federal legislation. That legislation will impact any US company doing business in Europe. So we can't overlook the fact that US companies are going to be regulated by virtue of what Europe is doing, and there's a lot of efforts within the TTC to try to harmonize. I think the key there is going to be standards and tools.

Then you look at the US landscape, and while we may not have new legislation, what we're seeing across the board and the agencies is applying existing laws to new technologies. The FTC is just one example of this, and some of the remedies that they have imposed have been discouragement of algorithms. They've just launched a big investigation into OpenAI. Is it an unfair and deceptive trade practice? The latest guidance that came out of OMB to the agencies was, "Invest a lot in trustworthy AI, including not only the development side, but the enforcement side."

So I would just ask you to come back to Scott's original question on what is the best path forward for the US in terms of unlocking the potential and mitigating the risk, taking from today that we are not in an unregulated world, and the regulation, the enforcement is intensifying, and what is the best pathway to achieve the goals?

Halie Craig:

I think that's for me. Yeah? Oh. Did you want-

Scott Wallsten:

[inaudible], then Hal.

Hal Varian:

I was just going to make a short statement. If you look at what Matt talked about this morning in terms of natural experiments, here's a natural experiment at a global level where you have a heavy regulation model in Europe and we think a lighter one in the US. So come back to me in 10 years, and we'll see how this experiment turned out.

Scott Wallsten:

We're going to bring Peter Brown back, and we'll talk about it. I know where you are.

Hal Varian:

That's okay.

Halie Craig:

No, that's a very good point. When I talk about regulation, I think I've talked about it a little bit broadly, and I should have clarified that what I was alluding to is more of the proposals, the new proposals we're seeing from Congress that I think are fairly heavy-handed, like, for example, licensing of certain AI models before they can go to market. I think mandatory algorithmic impact assessments or impact assessments for AI prior to being able to launch that product are also quite burdensome because I think a lot of them are inspired by the NEPA process, which is a whole other issue. We saw how that got out of hand very quickly. We just had to waive it for chips, for example.

But no, you're right. Europe has been ahead on data privacy, and I think that's a good example. When I was at Facebook, this also came up. We were always racing to comply with whatever the new EU regulation was on content policy or on data privacy, and the US has fallen behind, so to speak, because we have not regulated as quickly in that space.

I think also part of this will come down to international cooperation. Right now, in my view, we don't have an administration that is necessarily going to Europe and advocating for our innovation sector, because the FTC is sending ... They have detailees in Europe right now that are telling the Europeans how to implement DMA and DSA, which objectively target American companies that were in part successful because of our lighter-touch regulation in the early 2000s of the internet.

So I think part of it is going to be figuring out the international piece, which is really up to the administration. Congress can probably do something to direct those negotiations, but even with IPEF, for example, digital trade was pulled from those negotiations. I thought that that was pretty non-controversial to just say we should have free cross-border data flows and no forced data localization, but we don't have that advocacy right now happening for our innovation sector abroad.

So that's one piece that's difficult, and then you spoke to how agencies right now are just using the regulatory authority that they have, in particular the FTC, but we're also seeing it at CFPB, DOJ, EEOC. They put out that joint statement. So I think the role for Congress there is partly an oversight one, like, "What are you doing? Do you have the authority to do this?" Of course, the FTC thinks they do, but we're skeptical of that. You'll probably see something coming from us on OpenAI soon because we've been following that case very closely. You can stuff anything into a consent decree these days.

So that's another issue. We saw it with the Meta case. So part of it's oversight, but also I do think there's probably some space for legislation that does provide guardrails on where regulation can actually happen and who has the authority to do it because right now it is a bit of a free for all between the agencies.

Scott Wallsten:

I think we're going to have to stop there.

Halie Craig:

Sorry.

Scott Wallsten:

No, no, that was great. There's so much more we can talk about. The good news is we're going to. Well, first thank the panel so much. I thought it was a really interesting conversation. Don't go anywhere because we next up have a fireside chat with Oren Etzioni by Shane Greenstein, and

I will start to introduce them as they come up. Also, don't go anywhere because we don't have a break yet. Well, you guys can go somewhere because we can sit here.