

Barefoot Innovation Podcast with Elizabeth Kelly, Special Assistant to the President for Economic Policy, National Economic Council, White House

***Note that transcripts may sometimes contain errors and that transcript timing notations do not match the posted podcast**

Jo Ann Barefoot: I have been so looking forward to today's show because we have a very special guest. We have Elizabeth Kelly, who is the special assistant to the President of the United States for Economic Policy in the White House. Elizabeth, welcome to the show.

Elizabeth Kelly: Thank you. It's a pleasure to be here.

Jo Ann Barefoot: I'm so happy to have you and I have known each other for quite a long time and you've done a lot of very interesting things. Let's just start by taking a moment to have you introduce yourself, talk a little bit about your background and then tell people about your role. And then we're going to talk about the special executive order that the White House has issued on artificial intelligence.

Elizabeth Kelly: Absolutely. So like many of your listeners, I'm a lawyer by training. Got my start in the Obama administration, working on retirement, consumer finance, things like the fiduciary rule, making sure that there were caps on loans, made some military service members, lots of good things like that. And then actually helped start a FinTech company called United Income, was about employee number 10 and built that company up over a number of years before we sold it to Capital One. And then came back into government service. My present role is special of the president of the White House where I'm responsible for leading our efforts around financial regulation, as well as lot of the domestic and economic impacts of technology, including artificial intelligence.

Jo Ann Barefoot: And this year the White House put out an executive order on artificial intelligence. We know the whole world has been getting more and more focused on that topic. Tell us sort of what the genesis is for the executive order, what you're trying to accomplish with it, and then we'll go deeper on the content.

Elizabeth Kelly: Sure. So I should first off say that this is not even a midpoint in our work, but a continuation of our work. I think this administration and this president has been focused on artificial intelligence for quite some time. We started with the AI Bill of Rights, which articulated bedrock principles for ensuring that AI systems are safe, effective, and transparent, and prioritize civil rights, equity, and privacy protections.

A few months later, the National Institute of Standard Technology published the AI risk management framework to guide AI developers and deployers in evaluating and managing AI's risks. And this summer and fall we received voluntary commitments from 15 leading developers, including Anthropic,

Google, Meta, Microsoft, and OpenAI to promote the safe, secure, and transparent development of AI technology.

So we started with principles and voluntary commitments, and we've moved to using every tool in our toolkit to make sure that we are protecting against the risks and harnessing the promise of artificial intelligence. Indeed, the president's directive was to pull every lever of the federal government's disposal and to lead with substance. And I think that's what this order does, and the fact that we did it so quickly, for government at least, really speaks to the urgency the president feels. So happy to talk more about it. That's a little bit of the backstory.

Jo Ann Barefoot: Yeah, that's great. And we will link in the show notes to the order itself and to the other items that you're referencing so people can look there. So what are the primary takeaways from it?

Elizabeth Kelly: So it is an admittedly lengthy executive order, I think there's 88 pages to be printed on legal paper. So I will not go through all of it, but encourage folks to check out the links. But it's really structured around eight fundamental principles which flow directly from the AI Bill of Rights and other bedrock documents.

The executive order directs the establishment of new standards for AI, safety and security, the protection of American's privacy, the advancement of equity and civil rights. The EO stands up for consumers and workers, promotes innovation and competition, advances American leadership around the world and so much more. In the interest of time, I'll just touch on a couple of those buckets that I think might be of interest to your listeners. One is around the safety and security of AI technology, and this is an area where the EO really takes sweeping action to protect against AI's potential risks. It directs the Department of Commerce to develop guidance and standards for testing the safety of AI models, and it requires developers the most powerful AI models to share their safety test results and other critical information with the U.S. government. It also directs guidance for clearly labeling and watermarking AI generated content, which we know is key again, across so many domains.

The second principle is innovation and competition. While the U.S. certainly leads the way in AI innovation, the executive order makes sure that we continue that lead. And by catalyzing research across the U.S., one of the things it does is pilot the national AI research resource, putting \$140 million towards providing AI researchers and students access to key resources and data, and expanded grants for AI research in vital areas like healthcare and climate change. It also aims to promote a robust and competitive AI ecosystem by providing small developers and entrepreneurs access to technical assistance and other resources. And in order to make sure that we have the talent we need, it directs the modernization stream on visa criteria interviews and reviews, so that we can expand the ability of highly skilled immigrants and non-immigrants with

expertise in critical areas to study, stay, and work in the U.S. So the next bucket that I touch on is around protecting workers.

President Biden was fond of saying that he's the most pro-union, pro-worker president, and so it's no surprise that the president mitigating AI's risk and leveraging its promise means making AI work for American workers. That's why the EO directs the Department of Labor to develop principles and employer best practices for employers to mitigate AI's harms and maximize its benefits for workers, including by making sure that workers have a voice in how AI is deployed in the workplace. This work will also address issues including labor standards, workplace equity, health, safety, data collections, freedom to organize, the risk of job disruption and more. In order to make sure there's real teeth attached to this, the EO also pushes federal agencies to find ways to make employer adoption of these best practices a condition for the receipt of some federal funding.

Another principle is around advancing civil rights. This has been a key area for the administration starting with the AI Bill of Rights as well as the executive order on algorithmic discrimination. But there's more that we can do and that's why the EO requires the clear guidance be provided to landlords, lenders, and federal contractors to keep AI algorithms from being used and exacerbate discrimination. It also directs DOJ and federal civil rights offices to develop best practices to ensure fairness throughout the criminal justice system in the use of AI.

Also want to touch on implications for privacy. We know that AI exacerbates the already serious risk to American's privacy that exists. It makes it easier to extract, re-identify, infer, and link together data about people. And AI's reliance on data heightens the already rate incentives for collecting it. That's why I think it's noteworthy that as part of the executive order rollout, the president called on Congress to pass bipartisan private legislation. We're reiterating the call he's made in prior towards the union and op-eds and other public statements, and it's why we're also seeing the CFPB and the FTC use the full extent of their authorities to protect American's privacy. The EO builds on this work by mandating evaluation of how agencies collect and use commercially available information data brokers, and directing stronger federal privacy guidance. It also prioritizes federal support for privacy-driven techniques, including private sector research into privacy-enhancing technologies.

There's a lot more content in the EO, and I don't want to take up too much time going through all of it, but I think it just shows the really comprehensive approach that we're taking to make sure that we're addressing not only the potential future risks, but the risks in the here and now, while trying to capture the extraordinary promise of this technology through investments in R&D, in our AI workforce, in our intellectual property, other things like that.

Jo Ann Barefoot: Yeah, that's a fantastic overview and so many meaty issues there. Pardon me. I know your background is in financial services, and an executive order by its nature is heavily focused on giving directives or mandates to the agencies that are part of the administration and finance some of our key agencies, our independent agencies, so they're not technically covered by all of it. But talk a little bit about how you are thinking about the issues of financial services specifically.

Elizabeth Kelly: Sure. So as you know so well, AI's effects and financial services are cross-cutting. I'd highlight several dynamics that I think are especially important for understanding AI's effects and the administration's work on this area. One is that we know AI can have significant effects on how lenders allocate credit. So consider the risks of bias or discrimination in lending. Without the right safeguards, AI can't amplify pre-existing biases. That's why the executive order directs guidance for housing lenders on avoiding unlawful discrimination through the use of AI to advertise housing loans and reiterates that landlords cannot use AI as a basis for discrimination.

But on the other hand, AI could help mitigate the risk of discrimination bias offering ways to mitigate these biases or remove the decision-making. And the EO recognizes that fact and to capture these benefits, it encourages the Federal Housing Finance Agency to evaluate regulated entities, underwriting models for bias and disparities, and to explore automated processes as ways to mitigate them. Another key area is fraud. AI has long been used to prevent fraud financial services by, for example, helping banks compliance teams, direct patterns and datasets that lead to fraudulent transactions or illicit financial activities.

And recent advances in generative AI are enabling banks to further improve these measures. At the same time, we know that AI creates new risks for the integrity of information, including financial documents, while increasing malign actors' ability to impersonate customer's voices, steal their information, or break into their accounts, scammers are now using voice cloning to impersonate relatives to try and convince people to send money, or get around voice verification systems and they gain access to accounts. That's part of why the President's EO directs Commerce developed guidance for clearly labeling and watermarking AI generated content. And I think really importantly, make sure the federal government leads by example, by requiring federal agencies to label or authenticate the content they produce or publish in hopes that more actors the private sector will follow suit. It's also why we've been working in this administration to help develop promising technical solutions to detect AI generated content.

And in the case of voice cloning scams, terminate a phone call early or warn the receiver while the call is in progress. Of course, AI is also having severe impacts on financial services firms, back-office and compliance functions as it is across the economy. We know that AI can help automate all sorts of manual tasks, including tasks involved with polling and processing compliance-related

information, back-office data management and more. It's going to bring real benefits, saving companies money and hopefully leading to lower prices for consumers and eliminating mundane tasks for workers. But even these seemingly mundane uses do come with risks. In financial services all of this data is incredibly sensitive, highly regulated, and so that you want to make sure as you're introducing new tools and technologies that you're having all the privacy protections in place that you need, and that becomes challenging as you're standing up any new technology.

It's also incredibly important as we're seeing AI embedded in chatbots, that we're making sure that people are getting good and accurate advice, and aren't acting on information that might be incorrect in terms of a bank balance, investment advice, whatever such thing. Which speaks to why you're seeing so much of a human in the loop to be a backstop here. The last piece I just want to touch on is around market stability. And this is something that you've seen director Chopra, director Gensler, and other regulators speak to a lot.

We know that algorithmic trading is of course one trend that introduces greater volatility to financial markets, but in addition, AI can introduce risk to financial institutions for infrastructure capacity to operate. It could exacerbate cybersecurity risks, raising concerns about collapse using critical financial sector infrastructure, which is why part of the EO and part of the work that Treasury is leading is so focused on making sure that we have resilient and safe critical infrastructure going forward. The other sort of market stability related pieces around market manipulation, we know that DeepFakes could be used for market manipulation. In fact, in May, stocks wobbled after a fake image of a reported explosion near the Pentagon went viral before officials very quickly clarified the photo was a fake. And this was very quickly nipped in the bud, didn't cause major issues, but you can imagine that as a sophistication of voice cloning, deep fakes, everything else rises, you could see more instances like this and we'll have to be on guard in the future.

Jo Ann Barefoot:

Yeah, absolutely. So I think a theme running through everything you're saying is the sort of double-edged sword nature of this. So it could make things much worse, it could also make things much better, and in some areas, technology may be causing a problem, but the solution could in part be technology. Also, I think privacy enhancing technologies are an example of that, and the same in bias as you said. So I'd love to go even deeper on the concerns about bias. There's so much risk that either the data being used or the techniques being used, and the so-called black box nature of the analysis could introduce bias or exacerbate biases that are already built into the system. And at the same time, there's a lot of reason to hope that if we have more data and smarter analysis, we can make financial services more inclusive and more fair. If you're thinking about how people in the financial world should be thinking about the bias issue, do you have particular thoughts on what are the keys to getting it right? So

Elizabeth Kelly: I think you're absolutely right, Jo Ann. This is a technology with tremendous both promise and peril. Our potential peril as the president has said, and it speaks to the double-edged sword, but also the huge opportunities. I think we've also seen the regulators very clearly state that there is no AI exception to the law on the books. We have a lot of strong laws protecting consumers against bias and want to make sure the black box nature of the technology is not used as a get out of jail free card. And certainly the regulators have been very strong in their statements on that front. A couple of things to double click on, one, I think it's just incredibly important that companies are aware of the potential for bias at the outset and thinking proactively about how do they mitigate bias throughout the lifecycle of an AI system, both in terms of the data they're using to train the models, the testing they're doing of the models, the outputs it's generating.

This is an area where we can really benefit from stage released, careful rollout, extensive testing, and an iterative process. And so I think it's incredibly important to have that front of mind. And I think this is part of what we're seeing more generally in the private sector, is just making sure that we have systemic verifiable ways of testing for issues in AI systems. If we can't measure the bias, we can't manage it. And that's why I think it's so important that we're seeing with the private sector and the government step up and do its part here. For example, NIST is working to develop guidelines and best practices aimed in part enabling deployers of AI models to address the risk of biases. You saw commitments announced by a number of healthcare companies last week to make sure that their systems are being tested for bias with who are being deployed in clinical or other settings. And I think there's a lot the financial services sector can learn from those things and indeed a lot that a lot of these actors are already doing to continue to build on.

Jo Ann Barefoot: Thank you. You mentioned the competitiveness issue and the opportunity for these new technologies to increase competition and make markets better. Say more about that.

Elizabeth Kelly: The competition is something that has been a key priority for the president, starting with the launch of the competition council and the competition EO back in March of 2021. And it's something he's especially emphasized in his tech accountability agenda. We know that when it comes to big tech market consolidation lets large incumbents shut out start-ups and small businesses, leading to less innovation, worse products, and sometimes higher prices for consumers. And I think there's the risk that AI can exacerbate a lot of these dynamics across the AI value chain. High barriers to entry threaten to shut small players out from the market entirely. We know that there's huge cost involved in developing large language models. Think about the cost of the data, the cost of the compute, amount of cloud services needed, all of which means that most recent generation and foundation models can take over a hundred million to train with further models just increasing that cost.

And I think that there's a risk that it just means we're entrenching a couple of big players in ways that are negative for competition, for privacy, for a lot of the things that we care about. And it's why in the executive order you saw a real focus on maintaining a robust and competitive ecosystem. Part of that is by directing resources to smaller players through the national AI research resources I mentioned, providing dating computes to academics, entrepreneurs so they can participate in this ecosystem as well as making sure that our existing federal agencies are using their full scope of authority to promote competition in this marketplace.

You saw in the EO an encouragement of the Federal Trade Commission to use the full scope of its authorities, but you also saw an instruction to every executive agency to think about competition in their regulation of AI. The OMB MMO, which was a separate document released a couple of days after the executive order, the talk about how the federal government will be using AI even includes a paragraph directing each federal agency to consider in their procurement of AI what the effects and competition might be, and try and avoid and directing contracts in a way that could encourage or turn a blind eye towards self-preferencing for other anti-competitive practices.

Jo Ann Barefoot: That's so helpful. I'd like to actually turn to that last topic on use cases for the government itself. Talk about what seems most promising in terms of the opportunity to make government more effective, more efficient by leveraging AI.

Elizabeth Kelly: Well, this is one where I want your thoughts Jo Ann, because I know you've been working for many years on how we can use technology to improve government functioning, but I do think that it's important to note that federal government is already deploying AI to deliver better results for the American people. In 2023, federal agencies identified over 700 ways that they're using AI to advance their mission, improving health outcomes, addressing climate change, protecting against cyber threats. I'll give you a couple of examples, but for the real nerds among you, you can link to the spreadsheets on the OMB website that outline all 700 examples. For example, we've been using AI to predict infectious diseases and assist in preparing for potential pandemics, to anticipate and mitigate prescription drug shortages and supply chain issues, and to protect natural disasters, and preemptively prepare for recoveries, to provide timely notifications to keep people safe from severe weather events, to help forensic specialists detect anomalies and potential cyber threats.

And we know that AI can help us do even more than we're already using it for now. Admittedly, AI has its limitations and it's just that it is a tool and it may not always be the right tool, but when it is, it can be very meaningful. It can help expand agencies, improve their operations, deliver efficiencies and enhance the security of government systems. And that's why the AI executive order and the accompanying MMO take real steps to make sure that the U.S. government can lead by example in responsibly deploying AI. So part of what those documents do is reduce barriers to government's responsible use of AI, including barriers

related to IT infrastructure, inadequate data, and sharing of data, cybersecurity, workforce gaps. It helps agencies acquire specified AI products and services faster and more cheaply and more efficiently by creating vehicles for more rapid and efficient contracting.

And most importantly, you're seeing a whole of government talent surge to bring in new talent by leveraging existing programs like the U.S. Digital Service and the Presidential Innovation Fellows and using accepted and direct hiring authorities as needed, as well as upscaling existing employees by providing training for employees of all levels. We know just how incredibly important it's going to be for the U.S. government to lead by example in its own use of AI, both in terms of the example we set globally as well as in our ability to deliver the best practice services we can for our citizens.

Jo Ann Barefoot: I'm glad you mentioned the U.S. digital service. I had just finished reading Jennifer Pahlka's book, *Recoding America*, which I'll also link to in the show notes. I'm hoping to get her on the show. But the opportunity to put together these converging technologies, the digitization, and the AI is just so powerful. So you anticipated my next question, which was going to be for the government's own deployment, how do we need to think about safeguards and removing barriers? Do you have more thoughts on that or did you cover it?

Elizabeth Kelly: I guess what I would say is that I would separate the potential uses of AI into two different buckets, which is part of what we see in the executive order at the end of the MMO. We know that not all uses of AI are equally risky. There's some that are benign like autocorrecting text messages and noise-canceling headphones. And I think part of what you see in the EO and the MMO is that we recognize those distinctions and that we urge agencies to remove barriers, like I mentioned, to inadequate data or workforce gaps to the responsible use of AI, but put in place additional guardrails for AI systems that pose risks to the rights and safety of the public.

There's a presumption that certain activities are safety impacting if they're related to the functioning critical infrastructure like dams or electrical grids, involve emission of hazardous material, emissions or hazardous materials. And that they're right impacting if they relate to law enforcement, employment, government benefits, things that are so essential for people. And in those higher-risk contexts, the agencies are supposed to follow a set of protections putting in place things like AI impact assessments, real-world testing, independent evaluations with ongoing monitoring, public notification, consultation, equity assessments, and efforts to mitigate disparate impact, as well as making sure that you're using responsible data. And we're also trying to lead by example and live our values in our procurement and use of AI by promoting competition, and consulting with federal employees and unions when AI is deployed in the workplace. I think the goal here is really to make sure that we are all systems go on the use of AI in low-risk contexts while focusing

resources and attention on concrete terms, is this balance between protections and innovation that we're acutely aware of getting right.

Jo Ann Barefoot: Yeah. Let's widen the lens for a moment and look at the global developments. Obviously, every part of the world is grappling with these issues, and we've had a lot happening lately. We had the Bletchley Park Summit, the G7 Code of Principles, the EU Pass, the AI Act. How do you see these initiatives fitting together? Are there issues arising that raise concerns about global harmonization of how people are approaching these challenges?

Elizabeth Kelly: I think we've seen a remarkable amount of alignment. You can even speak to that from a timing perspective. It was the same week that we had the president's executive order, the OMB MMO, the Bletchley Park Summit focused on AI safety in London and the rollout of the G7 Principles and code of conduct for AI. And there's a lot of overlap between those documents. I think if you look at the G7 principles, if you look at the code of conduct, they're building on a lot of the work that the United States has done in the voluntary commitments we received from the 15 companies, as well as a lot of the principles and key priorities that we outlined in, first, the Bill of Rights, the MIS risk management framework, and then the executive order, which I just think speaks to our success in really leading with substance and helping shape the global conversation.

And I think we've benefited from the fact that there's been so much effort to really have robust international dialogues, both at the leaders level as well as between the regulator level, which has enabled global regulators to work very quickly in responding to new developments in AI, and to do it in a robust and aligned way. Now, obviously, this is the beginning, not the end of the process, but I'm really frightened by the progress we've made to date, and I think lots of folks would say the same.

Jo Ann Barefoot: Great. I know we're going to run short on time. Let me ask you to sort of step back and maybe share with us what excites you the most. We can see the potential here for some things to really change profoundly, not just incrementally, but really maybe solve some problems that we've never been able to solve before, or at least go a long way towards solutions. What do you think is most exciting?

Elizabeth Kelly: I think there's a huge number of incredibly exciting applications that could really help address a lot of the largest societal challenges of our time. For example, we're already seeing that AI has impressive abilities to help with drug discovery, narrowing down the pool of drugs that might be viable candidates. So we're able to move to testing and deployment much more quickly. In the climate space, we're also seeing huge potential. Everything from optimizing data center pooling, to helping lower emissions from air travel, improving micro-climate forecasting in ways that better enable renewable energy and carbon capture and storage. All of these uses that as we think about tackling disease, tackling

climate change, AI can be a really pivotal tool at our toolbox. And it's just the beginning of that process, which is why I think we believe it's so important to make sure that we're directing dollars towards R&D and really making sure that we're focused on AI for good and all the potential benefits it can bring along.

Jo Ann Barefoot: Fantastic. Last question. I can't remember when we've ever seen an issue that sort of exploded onto the scene so rapidly and is so hard to keep up with, like AI isn't new, it's whatever, seventy-some years old, but between the proliferation of data and the breakthroughs on generative AI and so on, all of a sudden everybody is trying to figure out how to keep up. Do you have advice on that? Do you have any favorite sources, book, podcasts, websites, anything you would recommend to people who are trying to keep pace?

Elizabeth Kelly: Sure. So there's obviously a lot of great resources out there. A couple I would throw out for up-to-the-minute industry developments, I confess to being an avid reader of The Information. I also really like the Import AI Newsletter for its thoughts into the research frontier, what the technical developments are in this space, I find that to be a good weekly read. And I make sure to listen to and read whatever's coming out of Stanford HAI and a lot of the other key academic institutions on this topic.

Jo Ann Barefoot: Right. So again, we will link to the executive order and the other resources that you're talking about in the show notes at regulationinnovation.org. But if people want to come and find the executive order and related materials, is there a best site for them to come to?

Elizabeth Kelly: AI.gov. Everything is there. I encourage folks to the head that way.

Jo Ann Barefoot: Fantastic. Elizabeth Kelly, you've given us so much to think about today. I can't thank you enough for being on the show today.

Elizabeth Kelly: Always a pleasure to chat with you, Jo Ann. Thanks so much for having me.