COMMODITY FUTURES TRADING COMMISSION

TECHNOLOGY ADVISORY COMMITTEE (TAC)

"AI DAY"

1:06 p.m. to 4:36 p.m. EDT

Thursday, May 2, 2024

Three Lafayette Centre

1155 21st Street Northwest

Washington, D.C. 20581

PARTICIPANTS

ANTHONY BIAGIOLI, Designated Federal Officer
CAROLINE D. PHAM, Commissioner
CHRISTY GOLDSMITH ROMERO, Commissioner
SUMMER K. MERSINGER, Commissioner
CAROLE HOUSE, Chair
ARI REDBORD, Vice Chair
NIKOS ANDRIKOGIANNOPOULOS
DAN AWREY
CHRISTIAN CATALINI
TODD CONKLIN
JONAH CRANE
SUNIL CUTINHO
CANTRELL DUMAS
TIMOTHY GALLAGHER
MICHAEL GREENWALD
STANLEY GUZIK
JENNIFER ILKIW
BEN MILNE
JOHN PALMER
FRANCESCA ROSSI

PARTICIPANTS (continued)

E. GÜN SIRER
JUSTIN SLAUGHTER
TODD SMITH
STEVE SUPPAN
COREY THEN
NICOL TURNER LEE
MICHAEL WELLMAN
ADAM ZARAZINSKI
KIRSTEN WEGNER
ELHAM TABASSI
SUNAYNA TUTEJA
TED KAOUK

P R O C E E D I N G S

MR. BIAGIOLI: Good afternoon. I'm Tony Biagioli. As the TAC designated federal officer, it's my pleasure to call this meeting to order. Before we begin this afternoon's discussion, I'd like to turn to Commissioner Christy Goldsmith Romero, the TAC sponsor, for the welcome and opening remarks.

COMMISSIONER GOLDSMITH ROMERO: Thank you. I welcome all of us who are joining us today to AI Day with a special welcome to our guests who are speaking today, as well as the TAC members, who I thank for their public service and serving on the TAC.

Special recognition, of course, and thanks go out to TAC Chair Carole House and Vice Chair Ari Redbord. And, of course, I want to recognize and thank the TAC designated federal officer Tony Biagioli, assistant designated federal officers Ben Rankin and Drew Rogers, as well as Scott Lee, Yevgeny Shrago, and Zach Coplan in my office. And I also want to thank all the CFTC staff who worked

on tech and AI Day, including in particular Shivon Kershaw for her technical assistance on the report.

I've also invited to be with us today Ted Kaouk, who yesterday was named as the CFTC's first chief AI risk officer. Where are you, Ted? Raise your hand. Welcome. Welcome, Ted. So it's nice for him to be here and be talking about this on AI Day.

While it is well known that AI has been involved in financial markets for many years, AI Day is about gaining a better understanding from AI experts of the impact and implications of the evolution of AI, including generative AI. AI Day is also about how to ensure the safety, security, and trustworthiness of AI systems in this evolution. Earlier presentations in TAC on AI have focused on this evolution of AI. Much has been said about the great opportunities that the evolution in AI could bring to solve many of the world's toughest problems, or, as President Biden calls it in the executive order, "the promises of AI."

On the other side of the coin are concerns that have been raised about the potential negative implications of this evolution of AI, or, as President Biden calls it, "the perils of AI." As a way to manage those perils, the TAC has been examining the concept of responsible AI. Responsible AI includes principles of fairness, transparency, explainability of the AI's output, safety, and security.

The concept of responsible AI at its core embraces values by which governments and organizations can align their AI planning, design, and deployment. A critical component of ensuring responsible AI is governance, which is increasingly becoming foundational to organizations and the U.S. Government.

The TAC has sought to support the current efforts on AI from President Biden and the administration. To that end, the TAC has already had presentations from the White House on the executive order on AI, the White House Office of Science and Technology Policy on the blueprint for

an AI Bill of Rights, and the Department of Commerce.

Today, we're honored to hear from Elham Tabassi, associate director for emerging technologies, Information Technology Laboratory, and chief AI advisor at the National Institution of Standards and Technology, or NIST, which is in the Department of Commerce. She will present on NIST's AI Risk Management Framework released in early 2023. As we learned when NIST previously presented to the TAC on its Cyber Framework, NIST is the standard setter in the Federal Government for technology. After all, it's in their name.

Just this week, NIST released supplemental material for its AI framework in response to the E.O. In the announcement, U.S. Secretary of Commerce Gina Raimondo described how the Commerce Department has been working hard to research and develop guidance needed to safely harness the potential of AI while minimizing the risks associated with it. She talked about the tremendous progress the Department has made in a

short amount of time and said, quote, "We are continuing to support responsible innovation in AI and America's technological leadership," end quote.

As the TAC helps the CFTC understand the impact and implications of the evolution of AI for financial markets, a whole-of-government approach will enable the CFTC to leverage existing practices, principles, expertise, and resources of other federal agencies and harmonize our approach where possible.

Along with Associate Director Tabassi from NIST, we are honored to hear today from Sunayna Tuteja, chief innovation officer of the Federal Reserve Board. She will speak on the Fed's approach to responsible AI. I have had the privilege of working with Sunayna on several occasions.

We're also honored to hear from Todd Conklin, who serves as Treasury's chief AI officer and deputy assistant secretary, cyber. TAC members are very familiar with Todd, as he serves as a member of the TAC. As many of you know, Treasury just

released a report on AI and cybersecurity, a report that came out after Treasury's consultation with 42 financial institutions. As Todd will discuss, one of the results of those consultations was that financial institutions are leveraging NIST's AI Risk Management Framework.

I'm also excited to hear a presentation from the CEO of Modern Markets Initiative, Kirsten Wegner. I was so impressed by Kirsten's congressional testimony on the potential impact of AI on financial markets that I knew that the TAC would benefit from hearing from her.

As the TAC examines the impact and implications of AI evolution on financial markets, it must be at the forefront of that examination that U.S. financial markets are the safest, most trusted, and deepest markets in the world. Investor and customer confidence and trust in U.S. financial markets is paramount to the ability of markets to continue leading the world.

As a market regulator, the CFTC has a mission to promote the integrity, resiliency, and vibrancy

of markets. This mission is necessary to maintain the leadership of U.S. financial markets. The CFTC has an important responsibility to safeguard U.S. financial markets and the trust placed in them. As part of this core responsibility, in past meetings, the TAC has examined the potential for AI to harm markets and to undermine confidence and trust in markets.

Today, the TAC discusses what the CFTC and organizations can do to prevent that. That examination and other AI subjects are contained in a report by TAC's Subcommittee on Emerging and Evolving Technologies. I'm very excited to hear from the subcommittee chairs, who will be presenting to the TAC their report on responsible AI in financial markets for the TAC's consideration today.

Similar to Secretary Raimondo's remarks, let me say that the TAC has made tremendous progress in a short amount of time and is continuing to support responsible innovation in AI and America's technological leadership.

The subcommittee has worked a long time on this report, and I want to give them special recognition today. Chairs of the subcommittee are Dr. Nicol Turner Lee, senior fellow and director of Brookings Center for Technology Innovation; and Todd Smith, director of information systems at the National Futures Association. These two co-chairs have made substantial contributions and led their subcommittee in drafting a report that expertly addresses in a very thoughtful way the impact and implications of the evolution of AI in financial markets. These are thoughtful, iterative, and foundational recommendations geared towards more responsible AI systems with greater transparency and oversight to safeguard financial markets. Given the AI expertise, as well as market expertise of the subcommittee, the findings and recommendations in the report should be taken very seriously by financial services firms and regulators.

I'm very grateful for the work that went into this report, and please allow me to recognize by

name the subcommittee members that drafted the report with the co-chairs Dr. Turner Lee and Todd Smith. I failed to do that on the DeFi report to recognize all the subcommittee members, so let me remedy that for this report.

They are Cornell Law School professor Dan Awrey; Better Markets director of derivatives policy, Cantrell Dumas; Trail of Bits founder and CEO Dan Guido; IBM AI ethics global leader Dr. Francesca Rossi; Themis Trading co-founder and co-head of equity trading, Joe Saluzzi; Institute for Agriculture and Trade Policy senior policy analyst, Steve Suppan; Circle vice president and deputy general counsel, global policy, Corey Then; University of Michigan professor, computer science and engineering, Dr. Michael Wellman; U.S. Department of Treasury chief AI officer and deputy assistant secretary of cyber, Todd Conklin.

It is an impressive report for the TAC's consideration, and I thank all of the subcommittee members for their contributions and look forward to the discussion today, so thank you for being here.

MR. BIAGIOLI: Thank you, Commissioner Goldsmith Romero.

We will now hear prerecorded opening remarks from Commissioner Mersinger.

COMMISSIONER MERSINGER: Good afternoon. I'm sorry that I'm unable to join you today, but I want to thank Commissioner Goldsmith Romero for using the Technology Advisory Committee to lead the way in addressing such an important issue for our markets. I also want to thank TAC DFO Tony Biagioli for arranging for such a great panel of experts.

Today's discussions exemplify the approach the CFTC should take regarding artificial intelligence, laying the foundation with a better understanding of AI, while pursuing our important agency goal of promoting responsible innovation. Whether we're discussing AI or any other innovation, new technologies often present opportunities for better functioning and more efficient markets, but, unfortunately, they can also present opportunities for fraud, as well as risks for customers,

regulated entities, and the economy at large.

However, intervention before investigation is a dangerous recipe for regulatory overreach. Regardless of the technology, we must remain focused on regulatory outcomes versus regulation of a particular type of technology. Our regulations must remain technology-neutral. If we get caught up in regulating the underlying technology rather than the outcome from the use of the technology, we face creating rules that do not keep up with technological advances. This means we need to understand the risks and opportunities of new technologies like AI. That's where all of you come in.

My hope is that today's meeting and the TAC's work on AI will be an opportunity for the CFTC to gather more facts and increase our knowledge by leveraging the expertise of TAC members and the speakers here today. After we gather those facts, we can apply them to our existing rules and governing statute and consider whether, and if so, what changes are necessary to best safeguard our

customers and ensure market integrity, while at the same time fulfilling our statutory mandate to promote responsible innovation. I have no doubt that, through this approach, we will be able to foster better and stronger derivatives markets as a result of AI, just as the CFTC has done during other equally transformative developments throughout its 50-year history.

Thank you all for taking time out of your busy schedules to be here today. We appreciate all your work, and I can assure you it will play a very important role as the Commission continues to discuss the role of artificial intelligence in CFTC-regulated markets. Thank you.

MR. BIAGIOLI: Thank you, Commissioner Mersinger.

We will now hear prerecorded opening remarks from Commissioner Pham.

COMMISSIONER PHAM: Good morning. Thank you, Commissioner Goldsmith Romero, for convening this meeting of the Technology Advisory Committee. Over the past year, you and the TAC have really been

leading the way at the CFTC on the issue of artificial intelligence, and I look forward to the continuation of those efforts. I also want to recognize the TAC designated federal officer Anthony Biagioli, the alternate designated federal officers Ben Rankin and Drew Rogers, the CFTC staff, and everyone who worked to make today's meeting possible. And, of course, thank you to the committee members for taking time away from your day jobs to share your expertise and provide the CFTC with valuable insight into advancements in technology.

While AI has recently dominated the news and policy agendas around the world, especially with the development of generative AI, the financial services industry has been using various iterations of AI for years to improve automation and efficiency under the supervision of prudential and other regulators. At last July's TAC meeting, I spoke about the importance of utilizing existing risk governance frameworks and risk management disciplines to identify, measure, monitor, and

control emerging risks and new technologies. One example is operational risk management, which includes technology risk, cyber risk, and third-party risk. I believe that model risk management is key for AI risk governance, and that one way to approach appropriate safeguards for the deployment of AI and trading and markets is by looking to the existing approach to risks and controls in algorithmic trading.

I look forward to the presentation of the TAC's Emerging and Evolving Technologies Subcommittee report regarding responsible AI in financial markets and to the presentations by our colleagues at the Federal Reserve, U.S. Department of the Treasury, and the National Institute of Standards and Technology on their current AI initiatives. I'm pleased that our partners at the National Futures Association will be sharing their expertise with us, as well as leaders from the private sector.

Thank you, Commissioner Goldsmith Romero, and to the members for your service on this committee.

I look forward to today's discussion on artificial intelligence in regulated financial services.

MR. BIAGIOLI: Thank you, Commissioner Pham, and thank you all for your opening remarks.

Before beginning our first segment, there are a few logistical items that I've been asked to mention to the committee members. Please make sure that your microphone is on when you speak. This meeting is being simultaneously webcast, and it is important that your microphone is on so that the webcast audience can hear you. If you'd like to be recognized during the discussion and you are here in person, please change the position of your place card so that it's vertical on the table. And, Chair House or Vice Chair Redbord, I will recognize you and give you the floor.

If you are participating virtually and you'd like to be recognized during the discussion for a question or comment or need technical assistance, please message me within the Zoom chat. I will alert our chair, vice chair, and our technical support staff either that you'd like to speak or

that you need assistance, and we'll go from there.

Please identify yourself if you're virtual before speaking and signal when you are done speaking. And please speak directly into your phone or microphone for optimal audio quality on the webcast. Importantly, please unmute your Zoom video before you speak and then mute it after you speak because this is all being broadcast publicly.

Please only turn on your camera when you're engaging in discussion. Otherwise, you can remain off camera. And if you are disconnected from Zoom, please close your browser and enter Zoom again using the link previously provided for today's meeting.

Before we begin, last but not least, we'd like to do a roll call of the members participating virtually so we have your attendance on the record. I'm just going to go down the line looking through the Zoom, and I'll say your name if you could just unmute and confirm that you're here, and then we'll be ready to get started.

Francesca Rossi?

DR. ROSSI: Here.
MR. BIAGIOLI: Ben Milne?
MR. MILNE: Present.
MR. BIAGIOLI: Dan Awrey?
MR. AWREY: Present.
MR. BIAGIOLI: John Palmer?
MR. PALMER: Present.
MR. BIAGIOLI: Christian Catalini?
MR. CATALINI: Present.
MR. BIAGIOLI: Michael Wellman?
DR. WELLMAN: Present.
MR. BIAGIOLI: Jennifer Ilkiw?
MS. ILKIW: Present.
MR. BIAGIOLI: Adam Zarazinski?
MR. ZARAZINSKI: Here.
MR. BIAGIOLI: Steve Suppan?
DR. SUPPAN: Present.
MR. BIAGIOLI: Sunil Cutinho? Sunil Cutinho, if you are present on the line -- let the record reflect Sunil just messaged me in the Zoom chat that Sunil is present.
Are there any TAC members who are

participating virtually whose name I did not call?

[No response.]

MR. BIAGIOLI: Thank you. With that, I'll turn things over to TAC chair Carole House.

MS. HOUSE: Thank you so much, and thanks so much to Commissioner Goldsmith Romero and to the entire CFTC team for convening us and your leadership on this core issue, as well as to the leadership provided by our subcommittee chairs on this incredible topic that we're going to be discussing all day, and also to the whole TAC, just a wonderful group of experts across everything markets and finance and emerging tech.

I have been in a lot of discussions lately focusing on AI and convergence of a lot of these technologies that we're going to talk about, cyber, AI, blockchain, and a lot of the core issues that regulators constantly struggle with are constantly at the forefront, and they're certainly present in the discussion on AI, how to balance and respect and where there's mutual supporting and also tensions between freedom and security, between

privacy and accountability.

And at these areas, this is exactly where regulators have to constantly consider and weigh, what are the costs, and what are the benefits? What are the costs either in privacy and intrusion into people's lives and markets and institutions, but also what are the costs if we don't impose obligations that help us understand technologies and potential impacts and risks and don't get the critical visibility and the enforcement capability in order to do something about it if things go wrong.

These are all core issues at play, whether we're talking about blockchain or certainly in the AI space. And across all these issues and these pillars with these equities, there's three pillars that I'm also seeing raised in every discussion that we have at the TAC, which is information, infrastructure, and identity. Those things are core to the AI discussion. I think that they're also manifested really well in the report that our incredible subcommittee put together. And,

ultimately, regulators are now having to think about what are the right measures and controls, the benefits, the opportunities, and the risks related to getting our information and data and infrastructure and identity right, secure, and trustworthy to preserve the future of secure digital economies.

So as we continue the discussion today, I'm just absolutely thrilled to be here and to hear from the incredible experts, including a lot of former colleagues and good friends, who are deeply, deeply respected across this space and was honored to be part of the subcommittee, really every subcommittee at the TAC focusing not just on admiring problems, but also presenting meaningful solutions, and to a Commission that has also demonstrated that it's not interested in just admiring problems, but to certainly find the meaningful way forward to address them.

So then I'll transition into our first presentation, and it is my absolute pleasure to introduce our first speaker regarding artificial

intelligence, who is a good friend and deeply, deeply respected colleague Sunayna Tuteja, chief innovation officer at the Federal Reserve Board, who will speak to us about the Federal Reserve and AI specifically on advancing responsible innovation, Sunayna over to you.

MS. TUTEJA: Thank you, Carole. It's delightful to see you. I'm sorry I'm not there on the analog, though I'm a little bit intimidated about how big the screen is in the room and how far I should back away from the camera. So thank you, everybody, and thank you to the CFTC TAC Advisory Council for hosting this super timely AI Day discussion and really bringing together a diversity of perspectives and experiences on this topic.

You know, clearly AI, gen AI is having a Cambrian moment, and there are no shortage of shiny objects or spicy opinions. And, as was well articulated by the opening comments, there's a lot of excitement about the promise of this technology. But I'm happy to see that more and more it's a very balanced perspective, that as much as we should

really interrogate the potential of this technology, we also have to acknowledge that it is not a panacea or without perils, and we've got to also understand what the downsides are and how we ought to be actively managing and mitigating and understanding those risks as well.

I think it's incumbent on all of us, especially when we see this nascent technology make its way up the S curve with an unprecedented velocity, and the fact that the technology continues to shape shift and evolve and we're seeing a new innovation happen in this technology, not in the span of years and decades, but in a matter of, you know, weeks and days. If any of you spent any time on Hugging Face, it is both exciting, but also daunting to see the developments and the pace of the developments in this technology on a day-to-day, hour-by-hour basis. So I think it's incumbent on all of us to really flex that curiosity and make sure that we're leveling up as individuals and on behalf of our institutions. And this conversation today, I think, is an excellent

anchor.

For context, I am the chief innovation officer for the Federal Reserve System. As I tried to explain my role to my Canadian parents, the Central Bank of the United States is not at all central. We are a solar system of 13 entities, 12 reserve banks and the Board of Governors, and in this role, I have the awesome honor of working with all of our reserve banks and the Board of Governors in advancing purposeful innovation for the Fed.

As part of our overall innovation portfolio, we interrogate a range of frontier technologies, ranging from technologies whose utility is in the here and now, and technologies that may seem kind of far out like quantum and we may not know when it's Cambrian moment is, but it's incumbent on us that we're learning today and kind of understanding its value for tomorrow. So as part of that range of technologies that we actively interrogate, we apply a consistent heuristic of responsible innovation, and that really is anchored in, you know, let's learn by doing. Let's not think about

these new things in abstraction or be assumptions. Let's actually put hands on keys and learn by doing and understand the technology and its potential, but also its, you know, peril and how we might manage it. So that's an important component of, you know, how we think about responsible innovation.

The second component of our heuristic really is, you know, when there's something new coming at us, it is easy to get distracted. I am, you know, guilty as charged. And with innovation, sometimes it can give it a bad rap because it seems like, you know, you're chasing something and we don't know what the value accretion is, so a lot of our responsible innovation focus is anchored in are we looking at this new technology and new capability, whatever it might be, in the context of solving gnarly problems? Are we starting with a problem statement? Are we starting with a thesis and then thinking about how this technology might help us solve it and kind of give us a level of focus and discipline.

And then the third component of our heuristic really is, you know, as we learn by doing and kind of we focus on solving gnarly problems, are we designing meaningful optionality solutions that might help us level up the institution for the present, but also kind of that call option for the future.

So I wanted to share that heuristic a little bit with you because as I zone in on how we are, you know, advancing our learning and our experimentation around AI and gen AI that you'll kind of start to see a bit of a true line. As I mentioned, this technology kind of continues to evolve and change, and so we have to remain quite agile in our interrogation of it. And if there are any sci-fi nerds in this meeting today, we can all attest that we've read enough sci-fi to know that AI is often cast as completely dystopian or completely utopian when, in reality, that is not going to be the outcome. It'll be somewhere in the middle, but that outcome will be defined by the choices and the decisions that each of us make, and

I think that's why having that responsible innovation framework is important.

The other reason having a framework around responsible innovation is critical, especially in large institutions like ours and yours and others, is there is a level of risk averseness that's built into the DNA, and it's necessary, right? Break things and ask for forgiveness is not something that's kosher. However, you want to guard against apathy, that anytime you see something new, while we may be very good at assessing the risk of doing something new, we also equally have to assess what is the risk of not doing this and kind of looking at it from both dimensions. So I think that a responsible innovation framework helps us guard against apathy, but at the same time, it also puts the right guardrails to ensure that we don't have a random thousand flowers blooming that, you know, may create unnecessary cost and duplication and risk, so taking that very anchored and prudent and pragmatic approach is critical.

So how are we going about with our responsible

innovation framework, and how is it manifesting? We really break down the framework along the lines of four E's. So the first component of the four E's is education. We very much take the stance, as was expressed earlier, that education is paramount, and education must be evergreen. For example, we host ongoing roundtables with the industry so we can be learning from outside-in perspective. But for internal for all of our colleagues across the Federal Reserve System, six times a year or every other month, we host immersive education sessions that are about leveling the playing field and completely inclusive to everybody who works within the Federal, and we tackle different components of this technology.

Just a couple of weeks ago, we hosted a half-day immersive education event around the topic of "with great powers come great responsibilities," right? Because we had done previous sessions, which was really all about understanding the potential of this technology and how leaders in the industry across different categories were applying

this technology. So the last one was all about, okay, how do we also then think about the risks, right, and some of the ones that were mentioned earlier from an explainability perspective, bias ethics, you know, you name it, hallucinations, but then really getting tactical about what processes, policies, and protocols do we have in place today that help us address these risks, which may not be new, but may just be a little bit different magnitude, but then also understanding what new risks are being created with this next generation of this technology, that we do have a responsibility to create updated policies or processes or protocols. So that's an example of really tackling different topics and helping the entire institution build new muscle memory around themes that we may already know and understand, but also net new themes that we have to start to think about and build for the future. So that's around education.

Our second E is enablement. You know, one of the first things we did when ChatGPT had its, you

know, blow-out moment Thanksgiving 2022 and everybody was like, oh, what are we going to do? Are we going to use it? Are we going to let our folks use it, et cetera, et cetera? One of the things, you know, that we undertook at the Fed, and again, it's accessible to all of our colleagues kind of really democratizing the access, is we unlocked as many of the generative AI capabilities, but we did it in a guard-railed environment. We have an experimentation platform called Launchpad, which, again, is quarantined away from all production systems and has the right appropriate guardrails, but what we wanted to do was unlock these capabilities wherein as appropriate so anybody within our organization could really get that hands-on-keys experience, right? Because the next best idea of how to apply this technology is not going to come from the chief innovation officer. It's going to come from somebody who is closest to the end platform, the end product, the end process, who knows where the frictions are and how this technology might be the ideal solution.

So that's what we wanted to unlock is that learning, but that learning by doing and enabling the access to this technology, again, in a very guard-railed, quarantined, sandbox-type environment.

Part of enablement, you know, in addition to the tech stack, was also about, hey, do we have clear policies and processes for all of our colleagues? You know, do they know what the guardrails are about the can and can't do, recognizing that a lot of these components are not going to be static, that these policies and processes will need to be dynamic and evolve as we learn. But we wanted to make sure that there was that consistency and that coordinated approach so, intentionally or unintentionally, no one was going to get themselves in trouble. So that was a lot of the work that went into the enablement component, and again, doing it together as a franchise.

Then the authority of our responsible innovation framework really anchors on experimentation. So this kind of culminates the

education, as we're learning, the enablement, which is kind of having that tactile, tangible experience with the technology and the right policies and processes and guardrails that enable us to curate the right use cases. So we start with a problem statement or a thesis statement and then run it through a very coordinated experimentation process where we assess, is this technology viable? Is it feasible in terms of a solution for this problem statement? And could this be a solution that is scalable and secure enough? But then also is there a desirability, right? Does the end user or the business for whom we're experimenting have that desirability? And that's really important because, you know, you want to kind of break down that sometimes unnecessary barriers, but we, the technologists and kind of the business leaders -- it's like that meme of all the Spider Mans that are pointing at each other, you know, so we wanted to make sure experimentation was very much business-led and technology- and security-enabled. So that's kind of the approach, you know, we're taking

with our experimentation.

The other component I'll share with the experimentation is it's helping us really validate where is the real promise, and what is the depth of the promise of this technology? Because we get use cases where like this actually doesn't need gen AI, vanilla AI with a sprinkling of RPA will suffice. So we also want to caution against that gen AI becomes the solution for everything. But it's also helping us understand the magnitude of the risks. What are these net new risks? What are existing risks whose contours might be shifting so it helps us think about both the promise side and the peril side and really learn?

And as part of experimentation, we also ensure that from day one we bring along our BFFs in risk, compliance, legal, cybersecurity because we need their counsel and their expertise but not at the tail end of the experimentation. We've taken the approach that they're at the table even as we are designing and curating the use case, which we found to be a helpful tailwind. And also, they bring

that other possible approach that kind of helps us, you know, learn with a bias for action, which I think is important.

And then the fourth E of the responsible innovation framework -- and we're not there yet because we're still in experimentation -- is where we really see the viability, the feasibility, and there's the business desirability. Then how do we think about execution? You know, where does it make sense for path for production? And again, just want to underscore we're not there yet. We're very much focused on education and enablement and this very curated experimentation phase. But it's important to kind of be thinking about it because where we see the value manifest, we want to make sure that, you know, we've got that pathway and kind of have the building blocks in place.

So that's a little bit about our process. As we think about use cases, for experimentation, we are really starting to see the value, the sprouting value, both from some of our internal work, but also kind of the outside-in perspective. It's kind

of the notion of five C's. Yes, I'm a bit of a sucker for alliteration, so the four E's and the five C's.

So our use cases as we're interrogating them, number one is around concision, and it's a bit of a made-up word. It's a combination of concise and precision. I think you all can likely attest to it, where you are ingesting or managing terabytes of multimodal data that just is not humanly possible to kind of go through. Could this technology add value by consuming all of those datasets and then producing a TLDR? That is not just concise, but it's also accurate, it has precision, and then the human becomes the auditor or the editor to it.

The second area is around content creation. You know, if you have given a speech about a topic X number of times, could the model write the next draft for you? One of the things we've learned is there's value in something like that, where the content and the perspectives and opinions are static. If there's a lot of volatility in the

content, then it's not going to be as useful or helpful.

And then the third one is really around code transformation, and in it, you know, there's the arena of code generation. There's the notion of code migration. You know, institutions that have tech debt and, you know, code that's written in languages that you may not have anybody to manage today. How do you transform that code into new languages and kind of help bring down the tech debt while you're also generating new code? And, you know, I spent most of my career in the private sector as a product leader, and it was always the bane of my existence to have to prioritize my product roadmaps because I never had enough awesome engineers and developers. So I'm hoping if code generation does prove its mantle, that this could be that great augmentation so we can ship more features and functionality without having to prioritize our product roadmaps too much.

And then the fourth C is really around customer engagement, like this technology, again,

be a copilot to the human where, again, you're doing a lot of ask-and-answers with, again, consistent knowledge bases, but knowledge bases where there's a high fidelity in terms of the dataset.

And then the fifth one is really catalyzing operational efficiencies where you're managing, you know, large data centers or cloud architecture or networking sites or other kinds of operational capabilities within an organization. How can this help bring about a new of efficiencies. I will underscore, again, these are all, you know, in experiments and what we're learning from the outside, but our current thesis really is that this technology potentially has the ability -- I'm a pilot, so I'll explain it in like flying terms. It's for this technology to be your copilot, but it's nowhere near to kind of fly the plane on autopilot. So that's at least how you know we've been kind of learning and kind of our current thesis around this.

So I'll pause there. And, Anthony, you can

let me know. I'm happy to take some questions or other direction from you.

MS. HOUSE: Thank you so much, Sunayna, for your wonderful and incredibly like enthusiastic and illuminating presentation on the great work that you're driving over at the Fed. I was thrilled when you went on board that team and thrilled to hear what you've been up to for the last two years.

So at this time, I would like to open the floor to questions and comments from the TAC members. If you're in person -- great, Michael, if you'll start us off.

MR. GREENWALD: Thanks, Carole. Sunayna, wonderful presentation. Thank you so much. Very impressive to see how the Federal Reserve is approaching this topic.

My question to you is, how do you feel about the talent recruitment in this space right now? Regarding the education, the experimentation, where are you, would you say, on the ladder of talent? And are you getting enough from what you need? And then how are you learning from each other?

MS. TUTEJA: Hi, Michael. Thank you. Excellent question. I talked about some of the tailwinds. When I think about some of the headwinds, I'm going to stick with my pilot analogy here. When I think about some of the headwinds in advancing this work, not just at our institution, but I think any institution, you know, compute costs is definitely, you know, high on top of the list. Data fidelity is another one because I think we all know, if you don't have the right data, it doesn't matter what kind of magical AI you overlay on top of it, you're not going to get the results or the outcomes that you desire.

Talent, Michael, is the third one. The way we have been approaching our current experimentation as we're getting going is we created an in-house talent network. So we put out, you know, a call within the Federal Reserve System, and said, hey, here are some of the skill sets that we're looking for. You know, if you are somebody that has the mental model of learning by doing, you like solving gnarly problems, you know, this may be the thing

that you want to kind of come contribute to.

So what we did is, you know, we created a talent network that has data scientists, cloud engineers, you know, folks who are pretty proficient in AI and LLMs today, business owners, because what that talent network enables us to do is, again, level the playing field for experiment because, you know, that product manager who is closest to the end product platform, customer knows what the problem statement is, so when they submit their use case, they may go, you know what, I don't code, I don't know how to actually do this, but I am the product owner who knows the problem statement, and then we can match them up with kind of our internal roster from the talent network.

Because what we wanted to do was break down the barriers of somebody saying, oh, I have this use case, but I don't know what to do with it, or I'm not an engineer, and we really wanted to make sure it's a big tent, and it wasn't just the technologists that were excited about it, that, you know, everybody at the Fed was excited about it, so

kind of creating that talent network.

The other way we've tried to augment is kind of working with external partners, right, collaborating with academia, collaborating with the private sector, and really learning from them and co-creating with them, which has really been super useful.

In terms of the third thing, as we, you know, think about our talent pipeline for the future, we are actively having those discussions about, hey, you know, how do we think about bringing in the next generation of talent for these types of roles like prompt engineers, right? If you ask the free version of ChatGPT, which stopped updating its data, I think, in 2023, what a prompt engineer is, it's going to go, I don't know what you're talking about, because there was no such thing as a prompt engineer in 2023. And, today, it's one of the hottest jobs on LinkedIn, so we've got to also kind of be -- you know, to quote the Gretzky idiom, you've got to skate to where the puck is going. So those are some ways of how we are pulling internal

talent, augmenting with external experts in academia and the industry, but then also kind of thinking about our future bench strength.

MR. GREENWALD: I think if you have Spider Man in the job description, it will be very successful.

MS. HOUSE: Thank you.

Justin?

MR. SLAUGHTER: Thanks. Sunayna, great presentation. My question's about the sandbox. You called it Launchpad, I think. Is that correct?

MS. TUTEJA: Yes.

MR. SLAUGHTER: So kind of a two-part question. First, how did you develop this? Was this primarily done in-house or off of, I assume, in tandem with an offer of an LLM? I'd be curious which one.

And then the second one is, you mentioned, of course, the network of the central banks, 13 reserve banks. How much have you made this available to other government agencies, either in financial regulation or otherwise, rather than have them build their own?

MS. TUTEJA: Thank you, Justin. So we built the Launchpad. I've been at the Fed now just over two years. I know, equal parts surprising and scary. So we built that shortly after I joined the Fed very candidly because this was the first time this institution had the role of the chief innovation officer. And as I did my listening tour around the institution, one of the pain points that was expressed often is, yes, yes, I know I'm supposed to innovate, but I can't even get access to a cloud account without going through 89 approvals and all kinds of hoops.

So it became very evident to me that if I was going to preach the gospel of innovation, I first had to make sure that we were removing their friction points and really empowering them and enabling them with what they needed so they could innovate. And one of those things was how do we curate a one-stop shop that's one-click access to all other contemporary technologies that you, A, either just want to learn, or you have a thesis and you want to experiment, and we actually built it

in-house.

The cool thing about the Fed is when you try to solve a problem, the first thing I like to do is do an asset inventory because within my solar system of 13 planets, someone, somewhere probably has thought of a solution or built a micro-solution for their local need. So in this case, we had a few sandboxes around the system at different reserve banks and areas, so we kind of brought them together and say, listen, we like this in yours, we like that in yours, and we kind of aggregated it and used that as our substrate, and then we continued to add new capabilities to it.

You know, it is a cloud native, so it's connected to all of the cloud providers. So when generative AI capabilities started to roll up, first, we started to enable the capabilities that are built into each of the cloud platforms, and then, you know, we kind of take a case-by-case approach to unlocking some of the other LLMs as and where appropriate, and again, working with my BFFs in legal and compliance and risk and cybersecurity.

So that's kind of the approach we've taken, you know, so yes, built in-house.

We have leveraged the Launchpad for some soup tech innovation work that we've done with some of the other central banks in the past, but it's not something that, you know, we bought widely syndicated, but if that's something that's of interest, I'm happy to learn more and double click offline.

MR. SLAUGHTER: Thanks.

MS. HOUSE: Nicol?

MS. LEE: Can you hear me? Okay. Well, thank you so much for your presentation. As Michael mentioned, using a Spider Man reference is actually really cool.

The other thing I wanted to bring up, though, is to sort of date myself. Many, many years ago, the Federal Reserve had a big conversation on online bias, algorithmic bias, data quality, et cetera. And I'll never forget this particular hearing that I participated on with Lisa Rice of the National Fair Housing Alliance where the Fed

had a particular interest in ensuring data quality over autonomous systems.

I'd love to hear a little bit more about what you all are doing to identify and mitigate bias and how that conversation has pretty much expanded over the last few years as we've learned much more about AI and its application in financial settings. And thank you.

MS. TUTEJA: Thank you. Thank you for the question. I will attest, I'm not a policymaker. I am just a mere mortal innovator, so I will offer my perspectives, and these are Sunayna's perspectives. You know, in the context of the experimentation we're doing, Nicol, you are totally on point. You know, the data integrity, data fidelity, as I mentioned, is super critical, but that's why we put such a big emphasis right now on that notion of learn by doing, right?

There is a lot of literature and a lot of conversations about minimizing, managing a bias, but from a process perspective, so if I'm going to tell our engineers, hey, you've got to consume

these models or create models where there's no bias, what does that mean, bias according to whom? Like, what are the tollgates? Like, what is the litmus test, et cetera?

So that's why, as we are doing experimentation, one of the things I'm really proud of is we have technologists and lawyers and policymakers and product people all sitting side by side, right? Because each of them brings in their specialized expertise so that, you know, as we're interrogating this technology, we're starting to kind of see the early results of some of the work. The technologists might go, hurrah, hurray, this is great, this is the result I wanted, but then from a policy perspective or product perspective or legal perspective, they may have other dimensions that we have to take a look at that's equally important.

So I might answer it that way to say we absolutely recognize that this technology is not a panacea and not without its perils and a lot of the things you talked about, but how do we truly understand the magnitude of it? Because then we

can say, okay, so the process, the policies, the protocols that we have in place today, does that suffice in managing and mitigating this? And if not, how do we need to level up those policies, those protocols, you know, and processes, if there's a change in the dynamics of these risks.

MS. HOUSE: Thank you, Sunayna.

MS. TUTEJA: Sorry, I would just say --

MS. HOUSE: Go ahead.

MS. TUTEJA: -- you know, we are also learning by doing. I know everybody is learning by doing, so I would also welcome follow-on conversations and welcome what you're learning from your experiments and work, you know, and advice you may have that, you know, we should also consider because, you know, we're all kind of on the same journey here. And the more learning and sharing like this forum, this facilitating I think will be a tailwind to all of us collectively.

MS. HOUSE: Thank you, Sunayna. And I'll take a little bit of chair privilege to ask the last question for this session.

So much of what you talked about, you highlighted and underscored how many stakeholders, how complex this is. You mentioned functions across legal and regulatory and policy and tech and security, and that's just for you guys thinking about how at the Federal Reserve System to implement and use these incredible, innovative technologies and also mitigate those risks.

Every single regulated institution in the financial sector is having to account for the same things, and I'm curious, coming from a regulatory body and knowing that part of what you're dealing with is trying to upskill talent and understanding across the regulators, the enforcers, the people in counsel's office that are going to substantiate, you know, different enforcement actions in cases, and I'm just curious if you have any closing thoughts about how regulators should be thinking about upskilling those capabilities and how they could take these lessons from how you guys are considering what type of tech to integrate yourselves. How then should regulators and

enforcers be looking at financial institutions that are integrating it and ensuring that they're taking as thoughtful and deliberate an approach as you are?

MS. TUTEJA: Yes, and I think a lot of that, Carole, is, you know, everybody's on that same journey, and I think that's why what I might say is having come from the private sector and now being in the public sector, I think it's actually really cool that the public sector institutions are just as engaged with this technology. You know, think of all the folks that are speaking here today and all the traction that's happening, you know, across the agencies around the U.S. Government. I think that that's a little bit different than maybe some of the previous technologies where maybe the industry is moving at a certain clip, but the public sector is a little bit behind or not even in the game. And with this, you know, I feel bullish because everybody's on the learning journey together, which I think can be symbiotic.

You asked maybe something for folks to

consider as they all kind of level up, you know, maybe I'll share an anecdote that I find, you know, that gives me both optimism around this, but also keeps me curious and wanting to keep leveling up is, you know, as humans, when we think about the utility and the application of new technologies, we often think about it in the context of what we know today and what we deal with today, right? And that's just kind of how it goes.

You know, I was going to ask for a show of hands, but I can't see it, but if you haven't heard or read up, I might encourage you to look up the anecdote of eternal September. So it comes from the internet history. So, you know, oftentimes, the Cambrian moment for the internet is alluded to the mid '90s, but the reality is, the internet had been kind of simmering, you know, from DARPANET to ARPANET becoming the internet. But until the mid '90s, there was a, you know, a coordinated onboarding to the internet, and it happened every September because that's when a new crop of students started at universities, which at that

time were the gateway to the internet.

In the mid '90s there was a company you may have heard of called America Online that sent everybody those square black things we used to call floppy disks, and said, put it in your computer and welcome to the internet. First of all, let's acknowledge how naïve and adorable we were that we would stick things into our computer people sent us, which you should not do, never, ever, ever anymore.

But that 1994-1995 timeline became known as the eternal September because millions and millions and millions of normie Americans suddenly landed on this thing called the internet, and it kind of broke the internet. It created all kinds of chaos, but it also unleashed unfathomable innovation because this technology that otherwise was in the hands of the few in universities or tech labs, et cetera, suddenly was in the hands of everybody, and it unleashed new ideas and unfathomable innovation.

And I kind of use that analogy sometimes to anchor myself with what's happening today. You

know, this technology on AI, gen AI, as nascent as it is, it's not completely new, but it was in the hands of a few. And what ChatGPT did was, suddenly, billions and billions and billions of normies started using it, so, yes, it's creating all kinds of chaos. But I feel bullish and optimistic that it's also likely unleashing unfathomable innovation that we just can't even think about. And I think that's exciting, but that also, you know, we have to make sure as leaders within our institution that we are staying curious, and we're hoping our organizations and our institutions level up to the right responsible innovation framework.

MS. HOUSE: Thank you, Sunayna. I love closing a regulatory advisory body on chaos and innovation. Thank you so much, Sunayna. You're wonderful.

I'll turn it over to Ari.

MR. REDBORD: Carole, thank you so much.

And it is AI Day at the TAC, so we're going to stay on that topic. Our second speaker on AI is

Kirsten Wegner from the Modern Markets Initiative, CEO there, and really looking forward to the presentation and another robust discussion, so thank you so much for joining us.

MS. WEGNER: Fantastic. Thank you so much for having me. I've learned so much and got some even reading material on the history of the internet that I can do afterwards.

It's just such a pleasure to be here. Thank you so much, Commissioner Goldsmith Romero, for inviting me and just being able to share my thoughts on the past several decades of innovation in the algorithmic trading space.

So I'm Kirsten Wegner. I'm the CEO of Modern Markets Initiative. I'm speaking on my own behalf, though, today as an industry practitioner who has been working with algorithmic traders, some of the first electronic exchanges in the early 2000s, as well as some of the early online brokerage firms. So I've been working as a lawyer and a policy advisor, really a translator and early explainer of that technology.

So I think just echoing what Sunayna just said about kind of looking at the rate of change and how unfathomable that can be, I mean, if you just look back at this rate of change, it's sometimes hard to distinguish the trees from the forest here. So looking at this rate of change here, I have a graph which kind of illustrates it. I mean, the first AI systems were around in the 1950s with an automated kind of chess system if you will. And since that time, up to like 2010, the rate of innovation has basically doubled every two years at the rate of Moore's law. But something happened around 2010 that we kind of broke through Moore's law and now change is happening at this exponential rate every six months. And so we can't predict the future, but we can look at the past and some of the lessons learned from the past 20 years in automated trading to kind of inform our conversations.

Estimates vary widely about the size of the AI market. Some say it'll be $300 billion by 2026, and that's no small change. So I think the real question is, as the industry grows, how do we

ensure that we employ responsible innovation, kind of what that means, to harness the best upside of this technology, while also ensuring that we have the building blocks in place, as I'll discuss, privacy law, protection against algorithmic bias, and a level playing field for market participants so we can have broader access of this technology, including governance systems and really promoting competition of the U.S. firms in a global market.

So here, again, we see at the very far right of this graph, at this point, some AI systems are surpassing human kind of intelligence in some limited areas of speech recognition. That's just been in the past 5 or 10 years. Now, that's some, not all AI systems. Say AI systems are vastly clunky, you need really robust, high-quality datasets to kind of surpass the human mind, and we're just simply not there yet in a lot of capacities. So there are many questions here. It's about the policy, right? If you need such a high amount of data to train an AI system, what happens if we can never have that data in the first

place, right?

So the really high-level observations here are that, first, AI and automation can and should be a tool rather than a replacement for humans. And we just heard earlier about the idea of a copilot. I think that's a very apt analogy. I think no one wants to cede total control to an AI system, so we have to have that human element and the explainability element.

Second, you know, I think it's important to recognize that data is the raw material going into these AI systems, so the AI system is only as good as its inputs. And, as I'll discuss later, this opens up a huge can of worms on measuring data quality, ensuring that there's not bias. I could do a separate presentation on that.

And then I think, as we talk about AI, there's some benefit in having an equal vocabulary here. What does this nomenclature mean? What is AI? What is generative AI? I don't think that we have clear definitions, and I think whatever the definitions are should be, I think, flexible enough

that they can evolve over time. But with so many buzzwords, it's important to make sure we kind of have a similar frame of reference.

And then I think it's important to focus on governance. It's vital to look at governance of the AI models and protocols to make sure they're fair, transparent, safe, and secure. So this could involve establishing best practices internally, working collaboratively so individual industry players don't have to reinvent the wheel.

Accountability, the public should seek accountability in order to have trust in AI systems. To make sure there are no gaps in frameworks, I'll talk later about, you know, the gaps in the privacy laws right now. We have this state patchwork. I think there's a role that we can play in filling some of those gaps to make sure we have an even and uniform system there.

And lastly, I think I'd be remiss not to mention third-party service providers. A lot of individuals and organizations may want to get into the AI space but may not have the resources

internally, so I think we should really look at some of these third-party vendors, what those business relationships are, what the liability risk is, privacy, data risk, and really make sure we're including those third-party participants in the conversations as well.

So moving onto the next slide here.

So some of the use cases in automation that we can be very familiar with, it's been used for a while in the financial services industry, we have CFO tools for dynamic pricing of goods. We have speech recognition, like, for example, if you're a trader, you can only listen to one earnings call at a time. With AI, this is something that's been around for more than five years, you could listen to multiple earnings calls at once. So that's an example of kind of surpassing the human capacity for recognition.

Alternative data analysis, anyone who's watched some of the HBO shows have seen traders looking at shots of like warehouses, parking lots, it could be a supply chain of boats going up or,

you know, shipping portals. So there are many different forms of alternative data out there that also can be inputs in the financial services industry.

And then I'd say the majority of the use cases of AI and financial services are likely to be kind of the less glamorous and more boring administrative functions of, you know, processing paper, reporting, and just keeping track of books, and just automating away, kind of the very mundane kind of entry-level functionality in the financial services industry.

Now, if we move to the automated trading industry that I've been working with, on the next slide, there we go, so this is a great picture from the 1980s of floor brokers on the trading floors. And if any of you have watched Wall Street, the movie, that's what the trading floors used to look like pre-automation. I don't know how many of you have been to a trading floor more recently, but it looks like the picture on the right where you have electronic traders reading orders using algorithms.

So the workforce has shifted away from these humans on the floor to really coders programming the algorithmic, you know, formulas for trade execution. So there's been a net increase in future workforce, but a complete shifting of skills.

So the underlying technology for this market automation is really algorithms that are using dynamic pricing the same way if you're buying -- and this is a crude example -- if you're buying an airplane ticket, there might be dynamic pricing, right? I mean, that's really what's happening with the stocks here. They're dynamically pricing continuously at a very rapid speed, different underlying metrics, whether it's, you know, consumer sentiment, it could be fundamentals about the stocks, anything driving the value, and so you're getting real-time algorithms making that pricing more efficient.

What that means to end users on the next slide is that, you know, this is one of the clear benefits of market automation, and that use of AI-

like technology is that over the past 10 to 20 years, we've seen a huge reduction in trading costs, so the average cost for trading has gone down by more than 50 percent over a decade. And if you go even further back, if you were like trading stocks in the '90s, you might spend $6 to trade $100 of stock. Now, that's pennies on the dollar.

Now, what that means over a lifetime of investment is that this graph on the right shows you have a 30 percent higher return, and that's just from the incremental savings in costs. So that's a huge benefit to end users that I think you can see play out in other areas of automation as well.

Now, that benefit has not been without challenges, which I'll go into on the next slide. This is a huge amount of information I'm going to try to cover, and these are some of the real, I'd say, scandals of the past 20 years that have come in part from industry. You know, there has been a need to explain how the technology works to the public, and I think bridging that gap of knowledge

between the technologists and the coders and then having, you know, individuals who can communicate how that technology works.

So we've read, you know, books about "flash boys" and high-frequency traders. Those were really algorithmic traders at the end of the day, who did deliver a benefit to the public in terms of narrowed bid-ask spreads, but there always were some bad actors out there. So if you read about frontrunning or spoofing, those were bad actors who were using that same technology to do things that were illegal. So my member firms had, for many years, worked to deploy technology to use our own AI systems and our own technology to turn in those bad actors to FINRA and to make sure that they didn't fall victims to fraud and abuse.

But this highlights the broader question, right? When you're using any technology, it can be used for good or bad, and it's part of human nature. There's a bell-shaped curve of humanity. So we can't ignore that lesson learned that there were individuals using the technology for bad and

that industry participants and regulators had to work together to develop technology and regulations to detect and deter those bad actors.

Flash crashes, you know, when you have an algorithmic trading system, no technology is infallible to a glitch, whether it's a phone, a computer or any other technology. And so I think it's important, when you look in the broader context of AI, that there are methods of having an off switch, right, that you can hit a pause.

What did we do as an industry over the past decades of financial services? Well, we have circuit breakers. They were finally tested during the COVID pandemic multiple times in the span of 10 days when we had all this volatility and stocks were going way up, way down. I think that is a transferable item to consider in the AI context, right? How do you turn a computer system off? It's, again, tying in that item of the human element, the copilot. Is there a way to turn it off if there's an algo gone wild.

Third, again, is kind of this gold rush to

data. I call it liquid gold, but it's just the image of just the huge value of the data that the AI system is only as good as the data going into it. How have we seen that play out in 20 years of automated trading? Well, one must look only at decades of litigation now between brokers and exchanges on the SEC side, which I think can inform our thinking here of, you know, there were special different data feeds and faster data feeds and colocation. You know, obviously, the exchanges are monetizing what they have, the data, and that is what one expects a business to do. But many brokers felt that there were different unfair pricing schemes, and it raised many questions and litigation that has, I think, on some level is yet to be resolved.

So this battle over data, this broader question, you know, it was -- I think I wrote an op-ed about the war of Wall Street maybe last decade, right, but this is going to play out again potentially in AI, right? Who owns the data? Who has access to it? Do we have data autonomy? And

I'll go into that as well of ownership as consumers.

Nomenclature I touched on earlier, like, how do you define -- there is still no definition for high-frequency trading. I think I wrote a comment letter to ESMA, a European regulator, a number of years ago, and no one can agree on the definition. You know, it's really an automated trading system, algorithmic trading, but I think people have found that they want to define the technology over the individual using it, that that same technology can be deployed by many different types of market participants, so the definition may vary. So how does that kind of influence our thinking on AI? I think, you know, we should look at principle-based approach not being too overly narrow because the definitions and use cases may change over time. So just that flexibility that we're not brittle in our thinking.

And then the importance of information sharing between industry and regulators, what we're doing today is just fantastic. I think we have to

collaborate. You know, I spend a lot of my time talking to academics, looking at math finance, as well as, you know, coders, as well as policy experts, right? How do we find some common denominator of agreement in this very rapidly changing space? The only way we can do it is working together and finding common standards, so I think more of this type of AI Day is a great idea.

Next slide.

Okay. So I mentioned data being liquid gold, and I think many of us at the table have been aware of kind of the patchwork of state data laws over the past two decades, that we have, I think, about a third of the states that have data laws right now. There isn't a federal framework, and I think it's really important that we look at a federal data privacy framework as just a foundation for our AI policy.

So I was really pleased this last month when I saw Senators Cantwell and Rogers introduced another bipartisan, bicameral bill on data privacy. Someone at a table recently who works at a grocery

store asked, do you support this or oppose it? I said, you know, at some point, just having any bill that establishes clarity, even if it's not 100 percent perfect, is probably better than having nothing. It does preempt some of the state laws, but not others, so there's still a little bit of, I'd say, room for discussion there.

Also on the federal front, we have Senators Warner and Kennedy who introduced a bill, not on data, but on ensuring that kind of risk factors related to AI are addressed. And I think that was discussed at your last TAC meeting, so very important that we have that federal focus.

The overarching policy questions, I think, that are being looked at both at the federal and the state level are transparency of data laws, consumer rights, like do we have a consumer right to data, data anonymity versus confidentiality, right? Like, what is a privacy right like? Like, maybe we don't want our data out there for everyone, but the government might still need a way to trace what we're doing to have accountability.

There have been some proposals for data broker registries. This gets kind of at some of that third-party industry participant question, as well as the very important question of how do you measure data integrity? What kind of systems should companies have in place to measure their data to make sure when data is put into a profiling algorithm, that there isn't a disparate impact based on race, gender, or other protected class factors?

So that is a lot of information, but I'm looking forward to discussing that in the questions.

I'll move on next to the state patchwork. I think that this slide is a perfect illustration of something that has not been generated by AI. It's human-generated. It's a little bit sloppy, and it looks a little bit visually like what the state framework looks like right now. We've got a lot of different laws, a lot of different thresholds for disclosure of data breaches, a lot of different rights for consumers, but it's in a way exciting.

Like the Constitution created states as kind of the laboratories of democracy, and that's what's happening right now. So we have three states and one city that actually have transparency requirements so you have to disclose to consumers if you're using their data for an AI system, which is kind of cool. New York City is one of them, California, Illinois, and Maryland. So it might be an employer. If you're videotaping a job interview, and you're videotaping their demeanor or doing something that would create a data input, you have to disclose that and get their consent.

You have three states that have clear frameworks now enacted for protection from algorithmic discrimination, and these are really requirements for self-assessment tools. They can vary widely in context. Some of them are just generally geared at prohibiting like business participants, if you're an insurer, from algorithmically discriminating against different potential insureds. You have ones aimed at the criminal justice system, pretrial risk assessment.

Believe it or not, some people getting in or out of jail may be determined by an algorithm, and so that's really important to look at in that context.

But we can look kind of peripherally at all those different legal frameworks and see, is that successful in that state, like kind of check out what is working, what isn't, and integrate that into a federal framework.

And then, of course, the most like standard, most common privacy laws at the state level really are surrounding privacy like authorized use, reuse. Do you have the right to delete your data? Believe it or not, most of the states don't have these rights, so our data can be collected, whether we're, you know, plugging in our phone in a car rental company or using an email network. We don't know how our data is being used, and most of us don't have any clear rights, depending on what state you live in. So I think it's important for us to look at this as a foundational question on the future of AI. It's so obvious, but we sometimes don't bring it up enough at these

meetings, that this is, I think, just foundational in building our house of AI policy.

So next slide.

Okay. This is a really creative idea I mention just to kind of tie things together. And Congressman Bill Foster loves to talk about digital driver's license, digital ID, but I kind of love talking about this idea, too, because I think it kind of wraps a bow on a lot of these issues. So again, right now, identification is by state. We have our state driver's license. We do have a federal-issued passport for foreign travel, but at, you know, the local level, our IDs are by state, and there are some states with digital IDs as well.

Right now, we don't really have a digital federal ID. It's debatable if that's a good or a bad thing. There's really an inherent question of tradeoffs again of privacy, autonomy, of traceability. Actually, USAID, of all organizations, has some really interesting graphics on like what a well-deployed versus poorly designed federal ID system could look like. And I have on

this slide some of the, you know, key points that they've made, which I think are informative, right? If you have a really well-designed system, you could use that for financial inclusion, get more people participating in the banking system, know-your-customer compliance, institutional inefficiencies. And if you're the government, you could give benefits easily and get tax refunds more easily.

So there's a lot of efficiency, right, but then I also think if we have a poorly designed system, you know, depending on who the leadership is, if you trace everyone's finances to their ID, if you don't like what someone says on social media, could you turn off their bank account? You know, if you have a bad government, you know, this could be a big privacy intrusion.

So I think, again, it's a tradeoff. This is really in the realm of sci-fi right now, but I think it's something to seriously look at because if we do have a federal privacy statute, it may raise the question of how do we tie that together

with a single ID? And if it's state-run right now, I think we still have a new patchwork evolving.

So I think we're ready for the next slide.

Accountability, okay, well, this is, I think, maybe one of the most challenging questions, right? I think I mentioned earlier, you know, human nature, there are bad actors, good actors. There may be good people doing bad things, and vice versa, right? It's as old as, you know, kind of philosophy and human nature, right? But if there are bad actors using AI, right, whether they are market manipulators, spoofers, frontrunners, people cyber attacking, or maybe they're foreign governments, or it could be also fraud, impersonators, AI-generated content that creates misinformation, deepfakes, right? There's so many potential use cases that we may have not even conceived, right? What can we do about that?

So I think, again, meetings like this are super helpful to have the private sector and public sector working together. I think in the private sector, you know, companies are self-motivated to

detect against bad actors because they don't want to be victims of like spoofing or other financial like illegal activity. How do you motivate people to want to protect their customers? Well, we want to retain our customers. We want to make sure that there's trust inherent in the system, that people will feel comfortable using the technology, and if they're bad actors going in to those customer accounts, you'll lose your customer base basically. So I think the private sector is incented.

But we also need to make sure we have structures in place, I'd say, through best practices, through internal protocols, that you can create accountability, fairness, transparency, and information sharing like between the different participants. So with the stock exchanges, if there was a bad actor who hacked into an exchange, what many did is they'd share the lessons learned internally within industry. We wouldn't want to basically advertise how someone got in, but it was really important to have that communication, the same way -- and this may be a dark analogy -- if

there's an airplane crash, right, and something goes wrong, you want all the airlines gathering together and doing kind of a postmortem. Like what happened here? How did this happen, right?

So if you have a breach in the financial services context with an AI system or some bad actor, I think you want to make sure there's information sharing that's happening, that people aren't embarrassed to show a point of failure, and that they share that point of failure or susceptibility to other participants in a confidential way so other industry participants can receive the benefit of that lesson learned.

I think, within the regulatory context, I've read since the President's executive order on AI was enacted like six months ago, that 150 staff have been hired. That's fantastic. I think we want the workforce capacity, I think we want the technology capacity to be able to keep up with that evolving space and to make sure we are leveraging the expertise of the private sector in informing decision-making.

So I think next slide.

So I think that is it, but I'd be happy to chat about any of the many topics I have brought up today.

MR. REDBORD: We're totally going to chat about the many topics --

MS. WEGNER: Please.

MR. REDBORD: -- you brought up today.

I'm going to open the floor up to questions and conversation for TAC members. I'll sort of just start with this. One thing that really struck me is we have -- everyone knows I like to talk about crypto and blockchain all the time, so --

MR. REDBORD: -- we have a DeFi subcommittee, and it really just struck me that so many of these issues, right, that you were talking about overlap. Like I could have closed my eyes, and you could have been talking about blockchain technology.

MS. WEGNER: Yes.

MR. REDBORD: And it's, you know, like a patchwork of state laws, a lack of federal clarity, market automation, really all these topics. More

sort of a comment, but if you have --

MS. WEGNER: Yes.

MR. REDBORD: -- sort of a reaction to that, I would love to hear your views.

MS. WEGNER: Yes, I mean, I think DeFi plays a really valuable role when you look at -- like, you know, one thing the financial sector has struggled with at times is like single point of failure, right? DeFi is so decentralized, you don't have like one single point of vulnerability.

I think also that if we look at the broader question of kind of looking at tracing the next step of generative AI and AI-generated content versus human-generated content, that DeFi and kind of decentralized finance has a place in that. And you can indulge me in this, but you could have a watermark to note the AI-generated content, and you could save that in a blockchain and have a traceability to identify, you know, is this real or not? And this impacts so many things, whether it's, you know, just an image or like election integrity, right? Is what we're seeing real or

not? Is that voice real or not? So I think having a way to trace content and tie it in with a blockchain and maybe use decentralized mechanisms is really exciting. It's really, really exciting. Yes.

MR. REDBORD: Yes, thank you so much.

Corey?

MR. THEN: Kirsten, thank you so much, fascinating presentation. A lot of what you said resonated. I'm curious if you've thought a lot about if consumers have the ability to easily exercise a right to delete their data or opt out or what have you, whether it starts affecting the reliability of the underlying data, the usefulness of the data. Does it end up skewing the data if like a certain class of people are opting out? And like how to like reconcile that sort of like inherent tension?

MS. WEGNER: Absolutely. I mean, that's such a brilliant question because, you know, the quality of the data is determined, you know, having a vast scope of data, size of data, and then, you know, a

lot of people have talked about eliminating, you know, redundancy of data, but you're talking about the opposite, right, not having two sets of the same data to throw it off by like deleting part of it. So I think that's a really important question, right, to look at because we don't really have any use cases of data being deleted yet. And it definitely can impact the outcome, I think.

MR. THEN: Fair, fair.

MS. WEGNER: And then also, how do you trace -- I mean, this is one thing, too, I grapple with, right? Once you've given away your data, right now, we don't know if it's been given away, and we don't really have a means of tracing it. Maybe blockchain can be part of that in the future, but how do you delete it? What is the impact? And then also, some states have gone further to discuss a kind of monetization of your data. Like can you get a financial dividend on your data if you own your data over a lifetime. California Governor Newsom has discussed this idea. I think Senator Warner talked about it. Nothing's been introduced.

And I think in Minnesota this idea was also discussed, right?

And like, people have tried to model out what that value would be of our data. It's really hard to model out. Some people have said, oh, well, if you divide the number of users on Meta by the number of people using it, it could be like $7. Those numbers are really arbitrary. I don't think we're quite there yet with the data, but how do you track the data? How do you maybe monetize it? And if you delete it, like, are you getting a refund on the data? Are you paying them back for the data? I don't know. You know, it raises so many policy questions.

MR. THEN: Right. So you're saying there's room for work here and --

MS. WEGNER: Absolutely.

MR. THEN: -- we have jobs in the future?

MS. WEGNER: Yes, absolutely.

MR. THEN: All right.

MS. WEGNER: I mean, I think we could probably do a whole four-hour session just on data alone and

monetizing data. And is it like a commodity, right? People have written -- there have been many thought pieces on this. Is data a commodity? How do we philosophically conceptualize it? Some people may say data is personal. It's like an organ. You can't sell it. We shouldn't make this something that's incented to be monetized because we'll incent more harvesting of data or that it'll have an adverse effect on kind of humanity that if we think our data is valuable, we might spend even more time on our screens because we could make money giving our data away, right? Like where do we draw the line, and what are the tradeoffs here? Is this something from an ethics perspective we want to incent?

On the other hand, we are giving our data away. It is being monetized, so should, you know, consumers participate in the upside of this? These are all policy questions that are very much open and I think exciting to think about.

MR. REDBORD: This sounds like an entire other tech meeting, and it actually speaks to your liquid

gold --

MS. WEGNER: Liquid gold, yes.

MR. REDBORD: -- expression for data. Yes.

Justin?

MR. SLAUGHTER: Yes, I was going to say, first off, to your point, Ari, there is actually a deal between News Corp and the Polygon blockchain to verify information. They've got the verify chain. This is a key point, though, Kirsten, because, I mean, the data has already been monetized by everyone but the consumers.

MS. WEGNER: Yes.

MR. SLAUGHTER: What we've done instead is we give the people the right to trade their data basically forever, to access individual websites, and then giving basically no responsibility --

MS. WEGNER: Yes.

MR. SLAUGHTER: -- to the companies to keep it safe or to share that back with the users themselves.

But on privacy, I'm surprised you didn't mention, though, that while we don't have a U.S.

law federally, we have in Europe GDPR --

MS. WEGNER: Yes, absolutely.

MR. SLAUGHTER: -- and now, because of the lack of action by Congress over the last 15 years, unfortunately, most companies just comply with GDPR and states. So we're regulatory takers to the Europeans rather than actually set our own laws, which strikes me as very dangerous for finance --

MS. WEGNER: Yes.

MR. SLAUGHTER: -- because we have really, I think, got a lot riding on the idea that we set our own laws for the financial system. We don't take anyone else's. And to lose that on privacy is a very dangerous camel's nose.

MS. WEGNER: I think that's absolutely right, and I think, you know, we do have some existing privacy framework. We have the Children's Online Privacy Protection Act, which was from the early 2000s. And, you know, we've taken a bite at the apple a couple times. And, of course, different regulators have some already existing authority. And yes, major like international companies are

also adjusting to a European framework. Also, California has a privacy law that's very much analogous to the European framework versus Virginia has a more business kind of friendly framework.

So, I mean, the question is, how do we harmonize this? Is the U.S. going to be kind of a leader in this field? Are we going to follow, you know, other countries, too, because we can't find agreement internally? I hope the answer isn't that we're, you know, regulating by default or legislating by default because we can't find consensus. I have a lot of optimism that there are a lot of smart people in the room here and with the companies in both parties that we can hopefully find agreement.

MR. REDBORD: There are a lot of other questions, and this has been an awesome conversation, but I'm so sorry, we're going to keep it moving to our next speaker. But thank you, everyone, for the participation and for the conversation today. It was a great presentation, great discussion.

Our next speaker, staying on AI, Elham Tabassi from the National Institutes of Standards and Technology, really looking forward to the presentation and discussion. Thank you so much.

MS. TABASSI: Thank you so very much for including me in these wonderful conversations. Thank you very much, Commissioner Romero, for having me here and the whole CFTC committee. It's great to see many of our good friends.

Let me start by saying a few things about NIST, National Institute of Standards and Technology. We are a federal agency and nonregulatory agency under Department of Commerce. We have a very broad portfolio of research from working on blockchain, quantum, advanced communications, cybersecurity, to AI, and a longstanding tradition of cultivating trust in technology. And we do that by participating in developing a technical sound standard, advancing measurement science and standards that make technology more reliable, more robust, more secure, in other words, more trustworthy. And that's

exactly what we have been doing in the space of AI.

I'll be talking about some of our recent work starting with the AI Risk Management Framework. A little over a year ago, NIST released AI Risk Management Framework or, as we call it affectionately, AI RMF. It's a voluntary framework for managing the risk of AI in a flexible, structured, and measurable way, flexible in several means and dimensions, mostly to be able to keep up with the pace of technology, but also for the different organizations with different resources be able to get the right guidance for managing the risk of AI.

And structure, we talk about that attribute, but we spent a lot of time in the AI RMF basically, the first part of that, to bring the community on a shared understanding on many different definitions, what is AI systems, what we mean by risks, what constitutes trust, what are the trustworthy characteristics of AI systems? And the second part of that focuses on recommendations on how to achieve trustworthy, responsible AI systems.

The third attribute is very near and dear to many of us at NIST coming from a measurement science agency, measurable, and that goes back to the fact that if we cannot measure it, we cannot improve it. So in order to be able to improve the trustworthiness of AI systems, the safety, the security, the privacy, we need to know what they are. The first step is to know what to measure and then how to measure them. And AI RMF talks about this.

Overall, AI RMF takes a risk-based approach. It puts the emphasis on the outcomes that need to be achieved based on the risks involved and avoid being a prescriptive approach on exactly prescribing the steps that needs to be taken, although it provides as an informative annex how the objectives of the different guidance can be achieved.

It's also a rights-based approach. It puts the protections of the people and planet right at the center of this, and for that reason, it also takes a sociotechnical approach. And knowing and

recognizing that AI systems are more than just data, compute, and algorithm, they are complex interactions of data, compute, and algorithm with a human that operates the environment that they're operating in and the individuals and the communities that are being impacted by the use of technology.

And so as we heard, context matters. Of course, you can look at it and take it as a more general abstract technology that can be applied to many different technologies, but then we get into the conversations of the valuations and risk management context and use case matters.

It was also developed in very close collaborations with a broad sector of the community. In developing AI RMF, as we do with anything that we do at NIST, we reached very far and wide and make sure that we include the voices and the insights and input and experiences of the community that developed a technology. That usually encompasses the expertise of computer scientists, mathematicians, and statisticians.

But we also very early on understood that we also need to get input and engage the community that studies the impact of the technology. That becomes the psychologists, sociologists, cognitive scientists. And again, all of these things go back to the AI systems being a sociotechnical system, and data solutions are not just only purely technological.

I should say that much of the work that we do at NIST, it starts with a request for information or multiple rounds of draft for public comment, workshops. And, for the sake of transparency, we post all of the inputs that we receive in terms of responses to the requests for information or comments on our drafts on our website, including the recording of our workshops.

So let me go on and talk about what the AI RMF is. So the first part of the AI RMF, again, tries to bring everybody on a shared understanding of what we mean by concepts such as AI, risks, trustworthy. There has been, as mentioned, a lot of high-level principle-based documents. So the

first thing we try to do is to get those principles, bring them to the practices, and then from the practices to actionable implementation guidelines.

So in bringing the community of understanding, I just want to highlight a couple of the things. You mentioned the definition of AI systems. We were very cognizant that a lot of good work has been done, so we didn't want to start from the zero and scratch, so we leveraged a lot of the work that has been done.

For example, for definition of AI systems, we started with the definition of AI system in the OECD AI recommendation, and then worked with the community to improve on that definition. And, interestingly, the EU AI Act was mentioned. AI RMF was released on a Thursday, January 26. On Monday, through a State Department colleague, we got a call from, you know, the European colleague that they realized kind of the limitations of the definition of AI systems in their documents when they saw ours.

Another definition that I want to point out is the definition of risks. That was basically built on or maybe it's exact code of definition of risk in some of the OMB and NIST document that talks about risk as a composite measure of the likelihood of an event happening and the consequence of that event. The consequence can be positive, negative, or neutral. Risk as a kind of positive, negative, or neutral ultimately makes risk management about maximizing the benefits, maximizing the positive users, while minimizing the negative consequences, impacts, and harms. And I think that's really important to consider when we are doing risk management of AI systems and also considering the risk of not using AI and not only think about the abuses and misuses that can happen.

Another contribution of the AI RMF was, again, trying to -- in the hope of bringing everybody on a shared understanding is talk about what do we mean by trustworthy? What constitutes trust? What are the attributes and characteristics of a system to be trustworthy? This all gets to the measurable

attributes and helps the developer to know what type of attributes and characteristics they need to build in their systems and also the evaluator to know what to evaluate for.

In short, working with a community, we came up with the seven characteristics of valid and reliable, safe, secure and resilient, privacy-enhanced, interoperable and explainable, fair with harmful bias managed, and transparent and accountable.

AI RMF, again, with the help of the community, provided some sort of a definition. Again, the whole point is bringing some sort of alignment in the work of the community. And then use the seventh characteristics in its core part of the document that provides altogether 72 recommendations categorized in the four high-level functions of map, measure, manage, and govern for assessment and measurement and management of the risks.

In the very high level, the govern functions provide guidance on basically the processes and

procedures, the roles and responsibilities, all of the embedding that's needed for successful, effective, and efficient risk management. The map functions provide guidance on understanding the risk and impact, understanding the context of use, and trying to, you know, contextualize those trustworthiness characteristics within that context of use. Guidance on prioritizing these characteristics was the most important for that particular context of use.

The measure function focuses on guidelines on measuring these different aspects of the trustworthiness, but also emphasis that there's tradeoffs among all of these things. So, basically, zero error doesn't exist, zero risk doesn't exist. These risks are going to kind of intertwine each other, so having a holistic and complete understanding of the systems and how all of these risks may interact with each other becomes important.

Also put a lot of emphasis on regular monitoring, so AI systems are most of the time not

deterministic, so you cannot just, you know, compare to the previous, you know, software product. Just testing one and put it into production is not advised. So regular monitoring, while it is in production, and deployment is also important.

And the manage functions provide recommendations on basically coming up with the responses to the risk, again, understanding that zero risk doesn't exist, and sometimes accepting risk is the way to go.

I mentioned that AI RMF has been built on a lot of good work. I should also give a shout-out to SR 11-7 and a lot of good work that the financial community had done on this.

Ultimately -- and we talk about in AI RMF -- that we don't want it to be reduced to a checklist. So it's voluntary, it's flexible. Not all of the 72 recommendations are to be implemented. And, to that end, we put a lot of emphasis making sure that it doesn't reduce to a checklist, a compliance regime, and there's a lot of thinking that needs to

be done on understanding the context, understanding the risks involved in that context, and coming up with a good response in terms of prioritizations based on the resources that are available, but also the risk appetite.

Ultimately, what we want to go for is the culture of risk understanding and risk measurement, so instead of putting all of these kinds of safeguards as sort of policy measures or other different measures at the end of a system, try to get all of this thinking as early as possible, hopefully, as the design part, to strive for building the technologies that -- I stole it from somebody in government -- that it's easy to do the right thing, it's difficult to do the wrong thing, and it's easy to recover if it is abused or misused. So thinking about how to build -- and this goes back to the developers and designers and a lot of innovators -- that how we can really build in those trustworthy characteristics in the system and test them during the design, during the training, and not only past deployment.

The last thing I will say on the AI RMF is that we also had our eye on global alignment. We understand that the discussions of the AI risk management, AI governance is a global conversation. So at the time of the release of AI RMF, we also released a crosswalks that basically talked about how to do the mapping or crosswalk between the AI RMF and several high-level policy or technical documents at the time of AI RMF release. That includes some of the international standards, but also documents such as draft EU AI Act because, at that time, it was in the draft state.

Since then, we have produced more crosswalks from AI RMF, and some of them have been done by international partners with, for example, AI governance from Singapore or Japan, and Korea also wants to do similar things, and also more of the international standards.

We put also a roadmap out at the time of release of AI RMF, again, informed by the input from the community on the things that's needed to be done on operationalizations of the AI RMF, but

also taking it to the next step. We put a revision cycle on AI RMF because we know that making sure that, as technology moves and gets updated, that this document gets updated. We put a playbook out that, again, provides a lot of informative references, informative guidance on how to operationalize each of the 72 recommendations.

Fast forward to October, the executive order was mentioned. The executive order, I'm told, is one of the most comprehensive executive orders in the history of executive orders in number of pages, but at least for me, it's also in the number of the actions that it asked many different agencies to do. The executive order is really centered on our obligations to harness the power of good in AI while minimizing its negative risks.

NIST was tasked with several tasks all around measurement and evaluations, and, Monday, we released four documents for public release. One of them is a profile of AI RMF for generative AI. That document talks about risks that are unique to generative AI or uniquely exacerbated by generative

AI. So it lists 12 risks for generative AI and talks about them.

There is also a document as a companion resource to an older NIST document, secure software development, that talks about secure software development for AI and how to improve the security of the software in the AI space.

A third document is a landscape of synthetic content authentication and detections. What are the methods and techniques today, and what are the gaps and challenges in implementation of some of those today. We have a task later on to provide guidelines for synthetic content authentications to work with the OMB.

The fourth document is a global engagement plan for promoting AI standards, so all of the work that's needed to basically make sure that there is a global alignment on understanding AI risk and AI risk management.

This is a lot of work. We have our plate certainly full, but we are also very grateful for the work with the community, the collaboration with

the community, and also the establishment of the AI Safety Institute announced by Vice President Harris on November 1 or 2. You know, we look at it as a new and exciting resource that can help us deliver all of this, including the consortium that we launched by 222 members in February. The numbers have now increased to 260 and growing, and that gives us a convening space to work more closely with the community.

I'm going to stop here for any questions that you may have, and thank you again.

MR. REDBORD: Thank you so much for the presentation, and really looking forward to opening the floor up to TAC members.

One thing that really struck me is it was probably the most succinct definition of risk management, maybe the best that I've ever heard. You said risk management is about maximizing the benefit of technology while minimizing the risk. And we spent a lot of time on this committee in particular talking about how do we maximize the benefit, whereas people in risk management tend to

focus on just minimizing the risk part. That really spoke to me.

MS. TABASSI: Oh, thank you. Yes, that was very intentional because if you just get stuck on thinking about the negative consequences and harm, the bigger picture of, well, what's the risk or negative impact of not using the technology is going to be out of the conversation.

MR. REDBORD: And a really good intro eventually to Todd, who's going to speak next. I know his report does deal a lot with the benefits in addition to the risks.

Other folks? Oh, Carole.

MS. HOUSE: Thank you. Thank you so much for the incredible presentation. I'm a big fan girl of the great work going on over at NIST. I would love for your insights to double click a little bit on one of the key principles and areas that you talked about, areas around transparency and explainability. I loved that when it was emphasized in the AI Bill of Rights that we published out of the White House working with NIST

and a bunch of other agencies because, to me, you have to get that one right in order to figure out the others, right? Like if there's no transparency or explainability, you don't really have a meaningful way to be able to assess whether you're effectively mitigating bias, driving security, et cetera.

I also know from work as a regulator in other spaces that there would be calls often from industry asking for more standards and guidance on this thing when industry hadn't really coalesced around standardized best practices and that government doesn't create these things in a vacuum. They build it with industry and, normally, once industry has sort of coalesced around what those standardized best practices are. So I'm just curious if you could share some insights with us about that, around standards or guidelines specific to that, including how potentially regulators, as well as institutions in the financial space, especially, should be thinking about and contributing to standards and guideline efforts

that you or regulators may be considering relating to transparency and explainability specific to their sector and use of AI.

MS. TABASSI: Thanks for that question. So, again, going back to unpacking the terminology, we think that explainability, interpretability, and transparency as three distinct concepts and provide some definition because one of the questions that, you know, two, three years ago, everybody would say that, well, this large language models, these large models, they have millions and now billions of parameters. Do we need to know all of those and how all of these things work? So that becomes explainability and what happens under the hook and understanding the model. And there's, of course, you know, debuggers and developers that need to know that.

The interpretability becomes kind of understanding the context and answer the why, why did the AI systems give that answer, and that becomes important. You mentioned the financial sector. There's already laws on the books that ask

for if, for example, a loan application is not approved, what is the reason. So if an AI system is being used, there should be a reason for that.

And so we had a publication that came out in 2022 that talks about these differences, talks about the interpretability, meaningful interpretability, how much information, information to whom so that we can make it more actionable.

And then the transparency, as you said, is really fundamental. The transparency is basically documentation and reporting of all of the different actions and the different decisions that have been made through the whole process of the risk management. So if some system has been deemed nondiscriminatory, how they define nondiscriminatory, what type of testing that has been done.

And because of the importance of the transparency in the playbook for each of the subcategories, for each of those 72 recommendations, we have a specific session for transparency, how to actually report that. So a

lot more work is needed to be done in that space, and, as you said, a really important concept, so thank you for the question.

MR. REDBORD: Commissioner Goldsmith Romero?

COMMISSIONER GOLDSMITH ROMERO: Thank you so much. Today's presentations have been so fantastic. I think we could be here for like 10 hours to discuss all of these. And I really appreciate all of the participation of everyone on the TAC because we could just go and go and go on this, so thank you so much.

I'm going to go back to something Kirsten was talking about in her fabulous presentation about how AI has been used for decades, including in financial markets, but, as you know, in other places, too. And so if I can just ask a very high-level question of what was so different about the evolution of AI that caused NIST to, in January 2023, issue a specific AI Risk Management Framework, given that had been around for so long?

MS. TABASSI: Yes, thank you for that question. And yes, you know, I joined NIST in

1999. My first project was to develop an open vocabulary speaker recognition. The project was walk in the room and say, computer bring up my calendar and, from my voice, recognizes my calendar versus yours. So many of us have been working in this space for a very long time.

And I think the graph that you showed that something happened in 2010, and that was the ImageNet and the birth of the deep learning, you know, the explosion of the compute and data that allowed a lot of innovations to happen, and that just gets -- you know, the innovation cycle, you said that's six months. I heard it nine months.

And we have one appendix in AI RMF that talks about how AI systems differ from traditional software systems and why we need that. And data is certainly one of those. As this model becomes bigger, they can have emerging behaviors that we did not know at the very beginning. They are more powerful, and as now, it's also the scale and speed of the adoption.

So AI systems present unique risks that we

hadn't seen in software systems before. We talk about the data. We talked about the non-discriminant nature of them, that they can change. So the previous versions, you know, the ones that we have before the ImageNet, a lot of those features, a lot of those inputs were handcrafted. So we humans had a much better understanding that what we're getting to the input to the system and how the system works.

Fast forward to the deep neural nets, all of those features now are being kind of figured out by the network itself. So it adds another sets of n dimensions of opaqueness to the system that now it's doing something that we don't quite understand those functioning of the AI systems.

And then fast forward to large language model and ChatGPT and others, these are much bigger, and then they can exhibit behaviors that it was not quite understood during their development. So these were some of the reasons, but there are more. Yes.

MR. REDBORD: I think we have one more

question, Steve Suppan online if you are with us.

DR. SUPPAN: Yes. So, today, we've talked about data as being the liquid gold of AI, and I just wanted to suggest that, you know, we should also be thinking about at least two other raw materials of AI, one being energy use and the other being water use. And I'm just wondering whether NIST has studied water use and energy use in the development of AI because, as you know, some studies have begun to appear about that use.

MS. TABASSI: Thank you for that question. Before addressing the water and energy use, I just want to go back to the data that you said. So the other thing about the data is that when you have all these large data, all of the biases baked in the data gets into the system, so that's something that we didn't have with traditional software. And a lot of the privacy issues and the security issues in AI systems of vulnerability, you put a positive note on it, a stop sign, our human brain can understand what the AI system cannot. So those are all the different attack vectors to AI systems.

In terms of the energy and water, in AI RMF we talk about the impact on the climate as something that needs to be studied, and I understand that a lot more work is now around this. OECD had started a new sort of a -- now I think it's about one year or six months or so -- a new sort of a study group around impact of AI systems on climate. We don't have a real study that generates empirical data yet on that. Thank you.

MR. REDBORD: Thank you so much. Thank you for the presentation and the great conversation.

I'm going to turn things back over to Chair House to get to our next presentation.

MS. TABASSI: Thank you.

MS. HOUSE: Thank you. Thank you so much.

And for our final presentation before we move on to discussing the report, Todd Conklin, chief artificial intelligence officer and deputy assistant secretary for cyber at the U.S. Department of the Treasury will present regarding the Treasury Department's recent report, "Managing Artificial Intelligence: Specific Cybersecurity

Risks in the Financial Services Sector." Thank you, Todd.

MR. CONKLIN: Thank you, Carole. Thank you, Ari. Thank you, Commissioner Goldsmith Romero. Really a pleasure, and thank you for the opportunity to present on Treasury's report, which was one of the many tasks that came out of the presidential executive order last October.

So the Treasury report, it's fully public. It's publicly available, a 50-plus page report that on the Treasury website. If you have any interest in going deeper on the higher-level overview that I'm about to give today, I think we do have some slide decks that hopefully you could all see well if you're at home.

If we go to slide 3, I'll give a little bit of background on the report. As we already mentioned, it was part of the ask of the White House that the Treasury start to dive into the impacts and use of AI and cybersecurity vulnerabilities in particular across the financial sector.

So to construct this report, Treasury

conducted, starting in December of last year, 42 in-depth interviews with multiple sector participants, ranging from our largest financial institutions to our smallest, and then we also included tech companies, core providers, some data providers, and AML and sanctions companies. We also heavily consulted our FBIIC, federal agency partnership, which is the regulatory committee of which CFTC is a member in addition to our broader financial sector coordinating council apparatus that's attached to the FBIIC.

So next slide.

What did we unpack? So we started with an overview of the current and existing regulatory framework, and this section again was done in consultation with the federal regulatory financial agencies, and the drafting of this section was led by the Office of Comptroller of the Currency.

So, throughout the interview process, most firms emphasized how collaborative their engagements with regulators have been as they've begun to explore advanced AI use cases. This is in

part because U.S. financial sector regulatory agencies do not focus on the specific technology per se, but instead address the importance of risk management, governance, and controls for the use of technology and the business functions it supports.

As a result, financial institutions have largely been able to incorporate existing guidance into their use of AI. Some specific examples of this include -- and many are listed at a high level up top there -- but rules and standards for risk management and information security, supervisory guidance on model risk management, standards and expectations for technology risk management, effective data management and governance, consumer and investor protections, effective third-party risk management, appropriate risk management controls and supervisory procedures for broker-dealers with market access, and also expectations for insurance providers on how to govern the use of AI. And we do a deep dive on each of those elements within the report if you have an interest in focusing.

So if we go to the next slide.

What are the current long-term use cases for AI across the financial sector? So most financial institutions, as previously mentioned, have been leveraging some element of AI for a decade, decades in some cases, especially for core cyber defense infrastructure like endpoint protection, intrusion detection and prevention systems, data loss prevention, and other network appliances. Although not specifically stated in the Treasury report, many of the anecdotes that we heard from our institution was that you can't do modern cybersecurity in this day and age without the use of AI tools. So that, I think, was a profound comment that we reflected back from many sector participants.

So participants reported that the use of AI has allowed for the automation and improvement of function, supporting risk management, fraud prevention, operational risk, and labor-intensive fraud and cybersecurity-related tasks. Generative AI introduces additional possibilities for

improving the quality, scale, and cost-effectiveness of risk management and compliance functions for financial institutions.

So next slide.

A little bit more on the cybersecurity and fraud protection focus, so financial institutions are taking an overall cautious approach to generative AI adoption and have prioritized trialing low-risk, high-return use cases. Participants also reported their use of existing resources like NIST's AI Risk Management Framework to support their enterprise policy. And of the 42 sector participants, every single one referenced NIST in particular as being a guiding foundational document that they attempt to leverage.

So our discussions with financial institutions also revealed that the use of in-house versus third-party AI systems varied significantly by institutional size. Most participants, regardless of size, appear to be using third-party solutions for at least some parts of their cybersecurity and fraud infrastructure. However, there's a clear

capabilities gap between larger and smaller institutions. Larger institutions tend to leverage a mix of proprietary and commercial data while smaller institutions heavily rely on vendor data to support their AI use cases.

We heard a strong desire from interviewees that improved fraud data information-sharing across the financial sector would greatly improve the accuracy and efficacy of anti-fraud models, especially for firms that do not produce a lot of in-house fraud data that they can leverage to improve their own detection systems.

If we go to the next slide.

So I think one of the notes that we tried to really highlight and why we fused together the cybersecurity and fraud issues, even though the initial scoping of the report was really specific to cybersecurity, is that cybersecurity very much is a team sport mindset, right? And I think what we heard when we asked the fraud questions was a bit of a different mindset, right? So often, the fraud functions are very insular amongst particular

financial sector participants, and there's not a significant amount of sharing across the sector, and there's limited infrastructure to actually do so. Some are in the process of being developed, and that model really causes a bit of a mentality where, you know, one firm just has to be stronger than their partner, right, as opposed to cyber, where there's a daily sharing of intel from government to the sector and amongst the sector through formalized vehicles that are really extremely well established at this point.

So part of what we hypothesize in the report is how do we begin the process of bringing that team cyber mindset and intel-sharing mindset into the fraud landscape specifically to close the data gap that exists in the AI fraud space for our smallest institutions? Really, it is a foundational problem that we identified in the report that we're really going to try to hyper focus on in our next step.

So in the report we also did a heavy dive on the best practices for managing AI-specific

cybersecurity risks, and we do go very deep on each of these, so if you are in the sector and you do want a reference point, this is a good place to start. So throughout the interview process, participants share their approaches for managing AI-specific cyber risk, which, as noted previously, is based on existing regulations and guidance. Our report consolidates these approaches to highlight the best practices we heard from participants.

The first is situating AI within enterprise risk management programs. This will integrate AI-specific risk management within the existing best practices of managing operational risk through three layers of defense approach.

Next, financial institutions should develop and implement their own AI-specific frameworks to guide their usage and risk management of AI systems.

Third, financial institutions are integrating risk management functions to cover the range of possible risks posed by artificial intelligence systems, and our next recommendation -- and this I

think was pretty well treaded in the conversations earlier today -- is to evolve the chief data officer role and map the data supply chain. In order to effectively harness the benefits of artificial intelligence systems, organizations must have a proactive approach to data acquisition, curation, security, and monetization, and especially most acute for our smallest institutions, which we already alluded to the data gap on the fraud side, it's a broad data gap for all AI development for their own in-house models, and that exposes them to a new third-party element as well, right, as that third-party data access becomes really the primary incubator for a lot of their AI systems and programs.

So we'll go to the next slide.

The report also then detailed a series of cybersecurity-specific challenges and opportunities, so we did lay out really a roadmap for how we could begin to combat some of these challenges, and we do have a playbook that we just recently leveraged in the cloud construct.

So Treasury in the middle of last year launched what we call the cloud executive steering group, which consists of regulatory agency heads and some CEO participation from the financial sector and also involves experts from cloud service providers. And collectively, across both regulatory agencies, the sector participants and also the cloud companies, we've begun to focus on a shared and common lexicon for cloud development, which is a very similar gap to what we heard in the AI space. And I think the most profound statement that we had early on in the Treasury report is there is not one consistent definition for what AI is or means or how it's going to be used for the financial services sector. I think that's a pretty striking statement, given that everyone has been active in this space for a very long time, and everyone has plans to continue to be active. So we have to start with some very baselining definitions of every aspect of this space so that, as we expand the regulatory environment, we are all speaking the same language across the tech companies, the

regulatory community, and also the financial sector.

And then, of course, I already alluded to the growing capability gap, which, frankly, in the tech space isn't new between large and small institutions, but the new element is just the absence of data on the smaller side to even do modeling, which can be one of the most profound challenges that we're going to have to collectively come together to fix.

And then we did work heavily with our NIST partners. One of the feedbacks, while every sector participant did leverage the AI Risk Management Framework as a baseline, there are opportunities to make that even more operational, so we are going to work with NIST to bring the sector in to try to get some of the more operational teeth behind some of the updates to that, so we're excited for that future partnership.

So, additionally, I would say I think a very appropriate next step, given that the federal regulatory agencies partnered in the development

and release of the Treasury report, the release on March 27 would be this body's report on artificial intelligence, which I spent a lot of time. And the Treasury team that drafted the March 28 report also has done a significant amount of effort into this committee's report, so hopeful that we can have a positive outcome on our discussion on that report today.

If we go to the next slide. Yes, I think I ran through it. We can go one more slide.

Yes. And so just general catchall is that there is a lot happening across Treasury in the AI space, so we're only covering a very small subset which is specific to cybersecurity and frauds. I would be remiss if I didn't mention that there are obviously other issues related to consumer protection that Treasury has obviously keenly focused on, and we're going to continue to monitor. And also, as chief AI officer, I'm responsible for all internal departmental development use cases, of which we have upwards of 100 ongoing right now in various stages of maturity that we're excited

about.

The most public that we've been in some of those use cases has been in our own internal fraud application use for our Bureau of Fiscal Services, for which, over the last few months, we've been able to reduce fraud payment activity by hundreds of millions of dollars for payments going out of Treasury's Bureau of Fiscal Services leveraging new AI models that we deployed relatively recently.

So that being said, I'm happy to take questions or comments. Hopefully, I'm not over.

MS. HOUSE: Thank you so much, Todd, for that great presentation.

I would like to open the floor to questions and comments from the TAC members. Michael, I saw you first.

MR. GREENWALD: Thanks, Carole.

Todd, thank you for the report.

As you look globally, given the Commerce announcement between the U.S. and the U.K., are there particular international examples in financial ministries that you're seeing alignment

where you're sharing best practices, given the work you just did?

MR. CONKLIN: Yes, a great question. At 6:00 a.m. this morning we had our quarterly G7 cyber experts group check-in, which I co-chair with Duncan Mackinnon from Bank of England. On our workstream agenda for the second half of this year is AI, so we did a deep dive presentation this morning. That workstream is being led by our Federal Reserve counterparts on the U.S. side of the house, along with a team from U.S. Treasury. So we're still in the process of scoping what that work looks like. We're going to attempt to use the Treasury report as a launching-off point. And just to give a spirit of the conversations that are, of course, confidential that we have across the G7, but the spirit is we're trying to find a new lane that would be appropriate for the G7 to focus on, right? And I think that's something -- there's so much activity in this space, right, so how do you, one, distinguish the activity you take, and what does that look like? So that's something we're

trying to hone in on what that actually looks like for the G7 partnership.

MS. HOUSE: Thank you so much. Great question, Michael.

Corey?

MR. THEN: Thanks for the presentation, Todd.

So something you pointed out up there is something I've thought a lot about, which is not just AI, but technology in general exacerbating the difference between large and financial institutions. And I'm curious, without having had the opportunity to read the report yet, how far out you've mapped what the potential consequences of that are, how you can ameliorate that. Do you see a wave of consolidation coming because smaller institutions simply can't keep up? I'd just be interested in exploring that.

MR. CONKLIN: Yes, fantastic. And I do want to be careful because we did only interview 42 sector participants, drawing sweeping conclusions in the consolidation potential space, but what was profound to me was that when we did interview

smaller institutions, in some cases, the CEOs of those smaller institutions joined the conversation, which is fairly rare for us at Treasury, right, especially in the cyber and fraud context. It's not rare in other contexts, but for that level of attention, I think, was telling. And that subset did have some very significant concerns of immediacy in that regard, not even a decade arc, but a much shorter arc than that. So it's something that we are committed to try to figure out as soon as possible.

MS. HOUSE: Thank you. Are there any questions elsewhere from -- oh, great, Ari has got his board up.

MR. REDBORD: Thank you so much, Todd, for the presentation and the report, and it's been fun to discuss it with you over the last few weeks.

I think one thing that makes it really significant is the fact that it was a survey of financial institutions to really kind of dig into what they're seeing. Can you talk a little bit about, anecdotally even, sort of what AI-enabled

fraud looks like in the financial institutions? What are they actually seeing? What's happening day to day there that they were concerned about?

MR. CONKLIN: Yes, great question. And it was on the slides, but I didn't cover it on the address, but I think what's new -- and we asked every single participant what's new here, right? Because just specific to the cybersecurity space, for example, we heard a lot about, you know, email phishing in the news. And when we started to hone in on, you know, how much is that impacting you, it's not really a new threat, right? But the email phishing is more precise. That is being observed. Translation abilities are being upgraded, right, so they're more sophisticated, but there are mitigations in place for that, right? So not that we disregard that. It's just there's mitigations in place that still apply whether or not the entry point is, you know, being constructed with AI or, you know, someone fat-fingering it, right?

So what is new that really then touches on your question with the fraud space is the ease at

which the bad actor or the fraudster can mimic audio and video. And why is it challenging? It's challenging to a very specific subset right now of -- and I think we talked about this on your podcast. It's challenging, especially for the high-asset customers who in many cases have a very public profile.

And oftentimes, unfortunately, financial institutions will create more of a frictionless experience for that subset of customer and reduce some of the multifactor authentication tiers and things to ensure a speedier engagement for that kind of -- that's been the new target area that came out of the interview construct with several within that construct, their voices and videos had been exploited to gain access. And that came up quite a few different times as anecdotes. We will keep it anonymized as to which firms.

But that is another example, though, there are mitigations in place. It was just the business risk caused those mitigations to be reduced, and that's something we tried to highlight in the

report, you know, a reminder that, you know, let's make sure because of this ability to mimic voices and video, that those controls remain in place for your entire customer set.

MR. REDBORD: Thank you so much.

MS. HOUSE: Yes, thank you so much, Todd. And congratulations to you and the entire Treasury team for a really just fantastic report, so wonderful.

That concludes our presentations so far. We're now going to take a short break and reconvene at 3:25 for then a deeper discussion surrounding the report that the subcommittee put together. Thank you. See you at 3:25.

[Recess.]

MR. BIAGIOLI: If everyone could please be seated, we're going to get started. Okay. For the next stage of the proceedings, I'm going to turn it back over to our chair Carole House.

MS. HOUSE: Thank you, Tony. We are now ready to explore our final topic of the day, what we're all here to talk about. It's a wonderful penultimate conclusion to this fantastic AI Day,

consideration and discussion of the Emerging and Evolving Technologies Subcommittee report on responsible AI and financial markets. The report was transmitted to the TAC members on April 26, and the plan for today is now to have a robust discussion regarding the report and, ultimately, a vote by the TAC regarding whether to adopt it and its recommendations to the Commission.

We'll start with an overview from Nicol Turner Lee and Todd Smith, co-chairs of the subcommittee, regarding the thinking and approach behind the report.

Over to you, co-chairs.

MS. LEE: Good afternoon to the full committee. We're on in terms of the report that you all have seen.

First and foremost, before we get started, Todd and I would like to send our grateful thanks to the Commissioner. Commissioner, we thank you for your leadership for this TAC advisory. We thank you for your leadership on this issue. And if you all remember when we first sat in this room

and we sort of got to know each other, this was one of your major issues, in addition to DeFi. And there is a third one that will go next, but this is an issue of importance to you, so we hope today that the things that you hear will be a relevant conversation to the conversations we had earlier today.

We also want to thank our DFOs and staff for all of your work in supporting our subcommittee. Without your invaluable support, this report would not be where it is today with regards to managing it within the agency.

We also want to thank our co-chairs of this committee as well. Thank you for your constant support and just the way you've helped us to frame -- and Carole particularly being one of the first people to put out a report -- helping us to frame the expectation for the quality of information that we hope that we will deliver to you and Ari today.

The subcommittee we will talk about in just a moment, but to the full TAC for again being here and being a part of this process. So Todd and I

wanted to start there because, as public servants, we often don't get a good thanks, and we want to make sure you leave with a whole hearty thanks for all the work that you have done even in reading the report that we sent around.

And then I'll just conclude and turn it over to Todd to give his thanks as well to the full subcommittee itself. We will post your names up on the board, but again, it's just really important for the subcommittee to know that we're very appreciative of the time and effort that you all spent as well.

So I'll turn it to my co-chair, and then we'll jump into the presentation.

MR. SMITH: Thanks, Nicol.

So I wanted to echo my thanks to Commissioner Goldsmith Romero, the support, the guidance that you've provided and the team you've put together, not only with Tony, Scott, everyone, all the DFOs were terribly supportive, productive, and I consider them a little bit of my CFTC family now, so you're welcome.

I also wanted to thank the subcommittee. It wasn't easy, but it was always constructive, and I think that's how the process should be, and I think that's because of the people on the subcommittee. So I thank them for their effort and participation, and learning through this process with them was really a great experience, so thanks to the subcommittee.

And then I also want a special shout-out to Nicol. As co-chair, I think we made a really good team balancing one another out, supporting one another, and I'm proud of the report that we're going to speak to so.

MS. LEE: And mind you he did that as I was completing my first book. It comes out in August, so I give a lot of love to Todd for tracking me down as well.

So we're going to formally begin our presentation. I'll just need some help with the slides. As we present this report, which is entitled "The Responsible Artificial Intelligence in Financial Markets: Opportunities, Risks, and

Recommendations. We submit before all of you -- the Commissioner previously read out all of the names, but I think I want to read them again. Todd was like, well, we don't have the time, but I think it's worth acknowledging all the people who are part of this committee again. And I'll just mention you by first name since your titles and affiliations were already mentioned.

Thank you to Dan, Cantrell, Dan G., Carole, Ben, Dr. Rossi, Joe, Dr. Suppan, Corey, Dr. Wellman, and to Todd Conklin who came on in his role after seeing all the stuff that you do. We're especially appreciative of your support of the committee process.

And we also want to recognize the support that Todd and I were able to get from our respective teams just with some of the things that many of you, particularly those of you who are lawyers, know that it's just the most grueling part, which is to fact check citations. And so they were very helpful in getting us to the bottom line.

And I also want to extend to Ben, who was

another person who was part of the CFTC team that kept us honest as we were working through the report.

Next slide.

This report, just to preface it, was not an easy report. As a researcher myself, and as we've heard from the various presentations, not just today, but the one that we had in January, there's not a lot of data when it comes to this particular mission of the agency. And the stuff that is out there, as I was corrected on many occasions, tends to be more consumer-focused than company- and client-focused. So this was one of those challenges, I think, for many of us on this committee to really think about how do we blend technological background, research background, and market background in constructive ways.

From that, I think the report objectives, as you have read it, really was designed to provide a comprehensive overview of AI adoption and use by entities. As you've heard a lot today, the word flexible, I think, was really undergirding our work

because there really is no textbook definition right now. I'm seeing Kirsten -- you know, we started just like you, probably a little bit more behind just because of the various disciplines and organizations that we came from.

But one thing we did agree upon, which you all have seen references in the report, is our collective definition of CFTC-registered entities, and that is very vital because I think, for the most part, it helped us to stay in our lane when it came to the basic premise of where we were starting.

The other things that we did in the report that you'll see as a second objective was to offer definitions of responsible AI governance and technical terms and really how they apply to the entities which we are referencing, as well as global financial markets overall. It's no secret that these terms are loosely used and often without common definition, and so it was very important to this subcommittee to give some meat to the bones of these definitions in ways that it was relevant to

how the report was actually going to flow, but more importantly, for us to develop our space and understanding and commonality among those terms.

I've been in this space long enough to know that there are other terms that we didn't include, that many of us have heard over the years, but I think getting to the responsible AI piece was probably in line but where we're seeing the conversation go.

The third goal of the report objective was access the executive, legislative, and regulatory use policy frameworks and current global guidance. So it was really important to us, despite us not having a lot of record and research in the industry itself, but to give some guidance on what we've seen across the board. And you heard a lot of that today, which we'll try not to be as redundant, but a lot changed also during the writing of the report in terms of the OMB recent guidance coming out. It was more recent guidance than the guidance that we actually reference, and so we did our best to be, again, flexible in the drafting of the report to

account for what was actually current and where we might actually go.

Next slide.

Another key objective for the report was to propose a series of use case scenarios, both real and hypothetical, for the application of AI. And I think you heard today from the prior agency representatives, you know, some of the applications when it came to fraud detection. And when we actually heard from Dr. Wellman in January and he started talking about super manipulative chatbots, we were like, we need to get him on the committee, too, right? Because these concepts, particularly for the scope of our work, were really interesting.

One of the things I'd like to note that I think we did quite well in this report is to not put, you know, the finger on the exact use case so that there was, again, some growth. This report in many respects, I think, is the beginning of an evolving document, something that can be very iterative as the environment itself changes. But we did prepare in this report, as you all have

seen, just really nice use case scenarios. I want to give a big shout to Dan for helping us with that where we were starting to think about some of those areas where we might want the agency to start putting a little bit more attention on.

The report also explored risk management strategies of AI tools and discussed and described some possible paths forward. Again, without a lot of literature there, there are some really good risk management strategies that we thought the prose and the policies and practices around those could apply to this particular space.

So in my experience of doing AI work, which is for quite some time now, it's real interesting, I think, the work that we did here because a lot of folks shy away from, you know, being very prescriptive. And I think what we were able to do was to give some behavioral practice that felt comfortable within the industry in which we all work. And then we're also applying this at the CFTC.

And we concluded with some proposed

recommendations. And we did something that in all of my years of being on a variety of federal commissions, if I appeared that I have done this before, I have, for a lot of years. We actually gave some guidance on who are the people we're talking to as opposed to, on many of the committees that I sit, making it the burden of Congress to do something different. We actually gave the staff some additional work to think about in those recommendations, as well as some of the organizations under the guidance of the CFTC and what we heard today, our partner federal agencies.

So that was the gist of where we wanted to go with the report. In case anyone was wondering, but why didn't they go further? Why didn't they start here? We tried to create a starting point where I think most people could fit and understand this without getting too technically wonky, and at the same time, being too prescriptive in a very flexible environment right now.

Next slide.

With that, I think most of the general public

will appreciate, as well as the CFTC-regulated industries, that we gave you a glossary to make sure that the things that we talked about in the report you also understood. This is just a sample of that glossary. The subcommittee, I love this idea generally. We kept finding that we were talking about these things, and we decided that maybe we didn't know the same thing about these things that maybe the public also needed to know. So the glossary really sets the tone for this.

Obviously, some things that are critical, we put up here on this slide, defining artificial intelligence. As you see with the definition -- I'm not going to read all of them because they are in the report -- making sure people were really clear about AI safety because I think that is the main reason why agencies like this are exploring this, that there's a safety component to this, particularly a safety component when it comes to the critical infrastructure of financial services.

Foundational models have recently come into the language of AI. For those of you that are

trying to understand what that is, read the report, but for the most part, these are the advanced capabilities that are really being placed on many of these models and tools that we see today, and then obviously generative AI.

What I also think that we did quite well in the subcommittee report is to not make this just a conversation around generative AI. Many of us right now who sit around the table, this is the new thing that most people talk about, and this was an area where we're able to give some thought into how this will effectuate the way in which the agency, you know, either regulates or recommends going forward. So, again, I just want to bring reference to the fact that there is a glossary in here with more than four terms, but these are some of the foundational terms that we'll bring up in the presentation.

Next slide.

So what is responsible AI? I just want to start here and then turn it over to Todd to go a little bit more into the risk management portion of

this. We decided for the purposes of this presentation to take a 65-page report and make it very concise with some of the key components so that we can have more time for dialogue.

We clearly want to make sure that folks understand that the framework that undergirds the majority of research in this report are around these typical properties -- we also define those -- of responsible AI. What does that mean? I think we've heard it today. It has a lot to do with fairness, right? What is the tradeoff when it comes to looking at fair models, and how do we evaluate the extent to which we're being responsible in their design, deployment, oversight, accountability?

Robustness, what we heard today from the Federal Reserve chief innovation officer is that this requires a very robust understanding of these technologies and not always with just technical insight, that you really have to be able to evolve with the space. At least that's what I got from it.

Transparency, having transparency matters and particularly in these markets, and so our friend here Kirsten talked a lot about, you know, transparency when it comes to where these models are being deployed. As I've said, I was really struck by your slide of the former trading floor and what the trading floor looks like today and making sure not only companies that are affected and changing and evolving into those markets understand the value of transparency, but the people who were procuring these products from also understand that as well.

Explainability, which we also talk about in the report, is another element of responsible AI. Our dear colleague Francesca really pointed that out in the work that she's doing to make sure that we understand how these models are being made and what's going in and out of them. It's a real big jump from our previous dialogues in AI about the whole garbage in/garbage out. It's actually gotten more sophisticated when we talk about it.

And then, finally, privacy and what does that

look like, particularly for end users of trading systems? You know, where do they fit in the ecosystem? So I was very appreciative, Scott and I, about the conversation you had on privacy, one, because we don't have to go into much detail, but it does bring up, I think, an element that goes beyond consumer privacy to how do we look at customer privacy going forward.

So I will end here, but I shared this to just give you a good sense of what's in the report, what's not in the report, right, what we were able to cover, what we still think we'll need to explore, and we'll share that in the recommendations, but just to do some level-setting in case you're wondering, you know, how far we went with the depth of the content.

So with that, I'll pass it over to my great co-chair, Todd Smith.

MR. SMITH: Thank you.

So, yes, in terms of level-setting, I think it's good. Tony, if we can go to the next slide.

One of the things with risk management that we

wanted to do as a subcommittee is set a premise. What are we discussing? Why are we discussing it? And that helped us with the route forward. So the premise with risk management that we took is that there are certain widely identified and well-known AI risks already in the industry and government and academia. So it was with that premise that we proceeded with the analysis as a subcommittee.

So next slide.

So, first, we wanted to identify, okay, well, what are those widely identified risks to a certain degree, right? So we wanted to work in specific categories. These are high level in nature, of course, but I'll step us through it.

Firstly -- and a lot of this content has already been touched upon, so I'll go quickly through it, but lack of transparency, essentially not knowing or being able to explain an AI model decision-making process. And what's key there, too, in terms of that bullet is transparency and explainability go hand-in-hand for the most part, so we wanted to definitely identify that as a risk.

Poor data quality, I think, again, we're referencing, Kirsten, your findings a lot. As we all know from systems, whether that be AI-related or not, you're only as good as your data. I've had other conversations as well. Data is the foundation. So, obviously, if that data is poor, that's a risk.

And we mention here overfitting. Overfitting is the noise associated with data. So anytime you're feeding that into any type of system, specifically with AI models, let's say, overfitting can cause unintended consequences relative to what the model and the users of the model or the model output is expecting.

And then poisoning of data, as it sounds, that's a risk associated with poor data quality. Whether it's on accident or on purpose, data poisoning can occur, and that alters the intended consequences or expectations of AI models.

Keeping on the data theme, mishandling of data, especially with sensitive data as we've talked about, and again, whether on accident or on

purpose, that's a risk.

Fairness, fairness runs the gamut in terms of the whole AI lifecycle if you will, so from data generation to model features, it's important that we speak to that because the consequences can be a reproduction or a compounding of the biases that have been spoken to earlier.

And then lastly, model concentration. Model concentration is when widely deployed models are used across the board and utilize AI foundational models by -- they're generated by a small group of entities that are shared by a large audience, and there's model concentration risks that results in that.

So next slide, please.

Okay. So partial list of use cases, this is extremely partial in terms of the amount there, but why do we talk about use cases here? As part of the subcommittee process, we've identified those categories of AI risks. Now what we wanted to do as a subcommittee was, okay, let's identify some partial list use cases. And as we get into the

recommendations, you all will see there'll be a reconciliation recommendation involving a gap analysis, but that's how we came to this conclusion as a subcommittee is that while they're widely known AI risks, what is unknown? And then in order to get to some answer to that question, whatever it may be, we wanted to develop a partial list of use cases and structure that into the conversation we were having.

So again, I'll go through these quickly, from left to right, trading and investment, that involves investment research and algorithmic and high-frequency trading, certainly AI use cases in those processes.

Next one is customer communications, customer advice, customer service. What business processes do use AI in that in that realm? Marketing, customer acquisition, customer retention, certainly.

Next, risk management, margin model monitoring, and then collateral and liquidity optimization, definitely use cases there.

Regulatory compliance, your market surveillance, your trade practice surveillance, know-your-customer, anti-money laundering, capital margin calculation. We can go on and on and on, but those are some high-level example use cases in regulatory compliance.

And then, lastly, back-office operations, trade verification, trade reporting, certainly use cases there in terms of AI.

Next slide, please.

Our next step as a subcommittee was to a certain degree leverage and understand and -- using Nicol's word -- level-set. Where are we in this space? And there are existing requirements that in the risk management program areas of the regulations where we have AI risks are tended to. And I'll go through the slide here and explain what these are.

So for FCMs, specifically, CFTC regulation 1.11, pulling one of the examples risk management program requirements under that regulation, technological risks, certainly tends to AI risks.

Swap dealers, CFTC regulation 23.600, pulling from that risk management program under that regulation, credit and operational risks speak to what can be posed potentially by AI in terms of a risk.

And then, lastly, our other bucket, if you will, registered entities fall under CFTC core principles and system safeguards therein. And specifically in system safeguards, the controls that are in system safeguards speak to AI risk.

And then at the bottom of the slide there, I alluded to this earlier, this is one of the recommendations -- not to get ahead of us -- but we want to have the CFTC have an engagement to identify any potential gaps there within the existing Risk Management Framework.

Next slide, please.

Okay. So what do we do going forward as a subcommittee? And part of our recommendations, our conversation led to this. Adding to and creating a more intensive risk management program, in addition to the existing risk management programs that are out there, that to us posed challenges and concerns

that we listed here from managing risk management independence if you will and where and how to fit humans in, over, under, wherever, on the loop. Where does that go? That's a challenging concern. And then vendor risk, that's a concern there, too, in terms of taking that approach.

What we thought was better and provided opportunities and benefits was taking a governance-based approach to initial AI risk management efforts. That presented a lot of opportunities and benefits. And it goes to our earlier speaker from NIST, and it's in accordance with NIST's framework where we take a governance-based approach from a firm level and wrap it into those processes so we can identify potential AI risks that we haven't yet seen and loop it into that process, have it really be process-driven. And then certainly accordance with the NIST framework, map, measure, monitor, and follow that process.

Next slide.

Okay. Related risk management issues, I'll go through this quickly because it's been spoken about

or these different areas have been spoken about. But training data, the quality of data, the type of data, and what does it represent. Again, an earlier speaker spoke to this, but data is key, and so that's a risk management issue that we wanted to make sure we tended to. The placement of humans, HITL and HOTL, common phrases, common terms, we joke about it, but it's a risk management issue that needs to be tended to.

And then, lastly, the growth of the talent pipeline. Recent OMB guidance that Nicol mentioned, staff capacity building around AI expertise. And then budget implications, so right with talent, there's budget implications, so we wanted to make sure to mention that as well.

Oh, I'm sorry. Yes, the spreading of misinformation. Again, going to Professor Wellman and his TAC presentation last time, we definitely wanted to include that as part of the risk management issues that we saw.

DR. LEE: Perfect. So then I'll go to the next slide.

So what you've heard from Todd so far, I think, is a direct application to the sector in terms of what AI looks like, which, again, I think is an extraordinary contribution, not just to the agency, but to the industry at large. We've been looking for these real-life applications I think many of you can attest to.

But in the governance side, like Todd said, it all relates back to governance because what we wanted to do in the report is to not make sure it was just one specific type of industry that is regulated, but that there are governance frameworks that could actually apply across the board. I like to consider these the work that the agency has to do, the work that firms need to do, the work that we need to see the technology community do, and the work of government. So governance overlays all of those verticals.

And if you notice, the first type is really very much related to technical cadence, you know, quality of data, model training, being able to look at the system deployment in ways that are critical

of the technology itself. Todd has mentioned it. What are the corporate policies and regulations that are existing in both regulated and unregulated industries? Oftentimes, with AI we think it's just one industry that should be the center of how we have these conversations, but there are adjacent industries that feed into much of the work of AI providers right now or companies that are using AI that are not necessarily using it as a central surface.

Strategies for identifying and mitigating technical and sociotechnical risks, a lot of that is embedded in the NIST AI Risk Management Framework, but it's also important for companies. Even if you do not think that this applies to you, it does. There are technical, sociotechnical risks that seep into models all the time to disrupt a lot of the stabilities that come particularly in the financial market. You have a lot of vulnerabilities that can be teased out, and I think that's that point of misinformation.

Government regulation and intervention, it was

mentioned earlier today, but I can't help but to continue to say that this administration has done a job of putting together the piecemeal of getting to the White House executive order on AI, starting with the national blueprint on civil rights, going into the NIST Risk Management Framework, having the OMB guidance as part of the White House executive order, and having agencies finally having chief AI officers who actually can take this on so that it's not just talk from the government, but it's real interventions that we'd like to make. And, obviously, we all know that we're doing much of this work to make sure that we're the global leaders in AI.

And then, lastly, strict guidance and guidelines on types of technology. This fifth type, as you've noticed, has sort of crept over to the other side on an indent, but this is really important as well because not all technology is created equal. I often find myself as a person who works in this space sort of thinking about artificial intelligence, but really thinking about

autonomous decision-making tools and really capturing it there, and I think we did a good job doing that in the report.

Next slide.

Why is this important? And this was an important point that we wanted to point out for those who may be also watching and reading this report who are thinking themselves on what governance strategies apply to them. It's important to have policies and procedures, as Todd said, to map, measure, and manage risk, ensure individuals are responsible for the oversight of AI systems, have the appropriate qualifications and training. These are different systems, and it just can't be, you know, always the person that did this before to do this kind of work.

Qualify informed decision-making by having a diverse team. Some of that has to do with demographic diversity. Some of it has to do with expertise. I am so happy to hear Elham talk about sociologists like myself who work in this space, but having those very interdisciplinary cross-

sector, cross-industry people at the table to discuss these technologies, it's a high-level concern.

Next slide.

Have robust engagement with relevant internal and external stakeholders. Establish clear communication and feedback processes with actual knowledge of the development, foundational training, and adaptation of these systems. Again, just on the governance side -- and we'll go into the recommendations -- this applies to everybody who may be around this table or listening via webstream, really making sure that there are communication processes in place and that you pick up where you might have made a mistake because that will be very helpful.

And again, I think our subcommittee was very honored to engage in this conversation where there was very little research because this is really important to the global economy and how we have the same type of protected interests here in the United States.

And address risks arising from the involvement of outside service providers. They are not exempt in this space, and so third parties matter, and I think we heard a lot about that today.

Next slide.

So with that, I just want to advise that this report is written in a way that anybody can read it, right? It's not too technical where it just like excludes folks from being able to understand this. That's thanks to our diversity of our own committee, who had to ask a lot of questions about what this was and figure out a space to contribute. But it's also written, as we've said, to be flexible enough to allow the agency to build up that knowledge base of participants that actually engage with them.

So we're just going to read the recommendations. There's more detail within the report. And then we want to just go to further conversation from our committee before we open it up to the full committee.

First and foremost, the CFTC should host a

public roundtable discussion, and staff should directly engage in outreach with CFTC-registered entities to see guidance and gain additional insights into the business functions and types of AI technologies most prevalent within the sector.

So, in other words, Anthony, your work is not done. And, Scott, we are handing this to you to do some more work that is part of building this knowledge base.

Second recommendation, the CFTC should consider the definition and adoption of the AI Risk Management Framework for the sector in accordance with the guidelines and governance aspects of NIST and to assess the efficiency of AI models and potential to actually say customer harms as they apply to registered entities. This is important.

A lot of discussion we had on our committee is the applicability of the NIST Risk Management Framework. There were some folks that didn't even know it existed in the way it does. We heard today that it can be applied to a variety of contexts, and it's something that is easily applicable

because it's not necessarily forcing you to do everything that is on that checklist, as we heard from NIST themselves.

And so I'll give you the last three, and then we'll wrap it up.

MR. SMITH: I just wanted to echo one thing, Nicol, that you had mentioned, and, again, going back to the DeFi report, the phrase that rang true with me when that report was presented and I think was carried on as a theme is that this is the start of the conversation.

DR. LEE: Yes.

MR. SMITH: And so I just wanted to echo that thought.

So recommendation 3, CFTC should create an inventory of existing regulations related to AI in this sector and use it to develop the gap analysis I mentioned earlier, potential risks associated with AI systems. And we spoke to already what the purpose of the gap analysis would be.

Next slide, please.

Recommendation 4, the CFTC should strive to

gather and establish process to gain alignment of their AI policies and practices with other federal agencies, including SEC, Treasury, and other agencies interested in the financial stability of markets. It speaks for itself.

And lastly, recommendation 5, the CFTC should work toward engaging staff as both observers and potential participants in ongoing domestic and international dialogues around AI and, where possible, establish budget supplements to build the internal capacity of agency professionals around necessary technical expertise to support those agencies' endeavors in emerging and evolving technologies. And that was our last and fifth recommendation.

DR. LEE: So they're a mouthful, but they're all written in the report, and you can read them at your own pace.

And so with that, we'd like to thank the subcommittee for all their great work at this and then just open up the floor to a couple of subcommittee members and the full subcommittee

membership following. So, Francesca, I know you're online if you wanted to say, as our chief AI expert on the committee as well, any comments about the report that we may have missed?

DR. ROSSI: Well, I don't think you've missed anything, but I'm happy, to speak in supporting the report. So, first of all, really, that it was a great experience to work with all the subcommittee members and the two of you, especially for your leadership, but also the collaboration, as Todd said, with your very constructive cooperation with the other members.

I think the report does a good job. And one thing that Elham also said previously that was very important to give a good balance between trying to understand the benefits but also the risks and not just focus on the risks and not just focus on the benefits. And I think that this is a very important balance that every document that tries to analyze the use of a technology and how to support responsible innovation should have. So that's the first thing that I think is very, very good in the

report.

The second one is, of course, these definitions that we give that give some clarity about what we mean by certain concepts that are relevant in this space. And in particular, I liked this deep dive into the various forms of AI governance. Nowadays, many stakeholders understand the importance and the relevance of AI governance, but not everybody means the same thing because, as we see in the report, there are five kinds of AI governance that we have considered. It can be governance of the AI lifecycle. It can be corporate internal governance. It can be governance given by public policy, by regulations, by standards. So it's very important to spell out all these different forms of AI governance because they all need to be considered and adopted and coordinated so that they don't, you know, contradict in a redundant way on top of each other.

Another thing is really the focus also on the Risk Management Framework, the NIST one, and to understand how to adapt it to this financial market

sector, which I think it's a very good starting point to then work on more.

And then, in terms of the recommendations, I think the recommendations do a great job in covering all the aspects because, basically, the recommendations say, first, you need to understand the capabilities, the risks, and the current uses of the technology of AI in the financial domain. Then, once you've done that, you can understand how to adapt and use the Risk Management Framework. You should also look at what our current regulation already covers and what instead is still to be addressed in regulation in other ways.

And then also, the CFTC, you need to collaborate possibly with other agencies to help each other, and in doing that, you should probably support all these activities with the raise in the technical and sociotechnical especially expertise and knowledge within the Commission.

So I think that this covers all the dimensions that need to be really improved not just in the CFTC, but in general. You know, raise awareness,

have more knowledge, understand how to adapt the current systems and management systems to your domain, work with others, see what regulation is already covering or whether there is need to do more, but support all these activities will really increase maturity in the sociotechnical issues and knowledge around AI.

So, definitely, I enjoyed working with the members of the subcommittee and under the leadership of Nicol and Todd, and I hope that the other members of the TAC and the CFTC would like this report.

MR. SMITH: Thank you, Francesca.

Any other subcommittee members? Or, Todd, I think you were going to speak.

MR. CONKLIN: Yes, fantastic. No, I really appreciate -- I was a late add to this work, so a lot of it was underway, but I really appreciate the partnership and the thoroughness of the report, really impressive. And I really think, in tandem, it's a really nice complementary piece to the Treasury report, and I think it's something that I

hope other agencies also strive to do.

So I think the more we could put out there in terms of baselining where everything is with the both the positive and also the risk management view, I think, is critical. I think it's something we tried to really achieve in the Treasury report, for example, which just does a great job at as well is telling the story of the sector's use in a positive way where I think a lot of times maybe the financial sector broadly gets some negative feedback in terms of not being innovative.

And I think actually the body of evidence in the AI space alone is that the financial sector and the financial services sector is one of the most innovative sectors in the AI space, and we should give them credit and also learn from them in terms of the best practices that have been deployed and also being mindful of the risks still that we have to work on. So I think this balances all that.

DR. LEE: Yes, I was going to say this committee helped me because I was pretty dystopian to start, so it was actually good to be around some

positive people for change.

I want to open it up to the subcommittee membership in general if there's any comment that anyone would like to put on record before we proceed to the full TAC. Do we have one from the subcommittee? And our esteemed Commissioner if you want to jump in.

COMMISSIONER GOLDSMITH ROMERO: Yes. So let me just say one thing. In listing out all the subcommittee members that contributed to this report, in my opening remarks I left out our chair Carole House because I just assume everyone knows that she is contributing greatly to everything, but let me recognize her that I know she wants to speak about the report.

MS. HOUSE: Thank you so much. Thank you so much, Commissioner, and thank you so much for going over this incredible report and the leadership that you two provided.

I loved this report and this process. The group on the subcommittee, the most incredible diverse representation of expertise and

perspectives. And I love that this report focused very much on understanding and mapping the reality of risks with a candid and, I think, very truthful acknowledgement of the opportunities. It doesn't pretend to have every answer. It doesn't pretend that the financial sector hasn't used and adapted AI before. They haven't been using AI in compliance or in conduct of trading like they really haven't been competing in finance, but there is a need for a call to action to understand fully these risks.

And that's really what this is. This is a launch. This is a beginning, and the whole subcommittee working on this with your leadership has provided a framework for the CFTC and other policymakers to take and thinking about how to apply that to these critical technologies. So it understands the wide variance of implementations and gave a really great overview of the landscape of different policies that are so critical.

Again, it doesn't ignore the fact that every institution that deals in the U.S. probably deals

with other jurisdictions, too, so I think that it was very thoughtful that you guys outlined those things.

I really appreciate the fact that regulators are having to consider with this technology how to balance providing guidance early and give a north star to responsible development of this tech, but also prevent crushing innovation in preventing being able to adopt and leverage innovative technologies or hardcoding in certain cyclical implementations into regulations that would prevent ensuring that that north star exists in the future. So, again, the risk-based framework that was really recommended by all members of the subcommittee and as echoed in the report, I think, is incredibly wise and poignant for policymakers to consider.

Also, I just wanted to note more as something that I'm looking forward to as we take the recommendations, so it's more going forward about the report, the convergence of AI with the other areas that we're focusing on. I know I mentioned it a bit in my opening statement, but AI with

cybersecurity, in a very good way, right, like AI is enabling more real-time detection and understanding of cyber indicators to enable our ability to defend against more sophisticated cyber intrusions.

But it also, as consistent with the purpose of the executive order and coming up with this term of dual use foundation models, essentially, provides what could be argued to be a munition and human-level capacity for command and control of incredibly sophisticated malicious cyber operations in the next few years. So we need good AI to help us combat the malicious use of the bad AI.

I also think that with blockchain and AI, there are some really interesting convergences of those two technologies, reinforcing each other's best and worse traits in really fascinating and interesting ways. Smart contracts could be truly smart and self-optimizing, and what does that mean, depending on what that smart contract is up to? And AI helping to solve some of the scalability problems of blockchain, maybe blockchain helping to

solve some of the transparency and auditability problems of AI.

So some really interesting work that I'm looking forward to us continuing with the TAC across all three subcommittees since I'm sure all of us have been talking about AI no matter which subcommittee that we were in. But it was a wonderful report, and it was an honor to be a part of it.

DR. LEE: So I just want to put for the record, since you are both a subcommittee member, as well as our esteemed chair, and if there are no other people on the live screen from the subcommittee, that we surrender this portion over to the TAC.

MS. HOUSE: Sure, thank you. And then we've got two more subcommittee members.

DR. LEE: Yes.

MS. HOUSE: So, first, Corey, I think you had your tag up first, and we'll go to Timothy and anyone else that has comments.

MR. THEN: That was just up from before, but

something very dramatic there, right?

But I think we should give a special shout-out to Nicol and Todd because they labored earnestly and doggedly for many, many, many hours doing this, so thank you guys.

MS. HOUSE: Thank you, Corey.

Timothy and then Stanley.

MR. GALLAGHER: We all work in emerging tech. It's an absolutely tremendous report. I like the way you guys didn't try to boil the ocean, right? You set it out there, and a lot of your work will be built off in the future, and I hope to be part of building off in our group as well.

Just a quick question or recommendation, too, applying the RMF NIST framework, which, you know, I thought was a great idea, do you see that as being like a governmental function or an internal function for the company or like even like a third party like an external accounting firm type of function?

DR. LEE: So I think the jury's still out on that one. I think we're going to see a lot more of

a market movement towards the Risk Management Framework only because I think it's coming from the government agencies as a guidepost for what we want general corporate governance to look like. But I think, as it evolves, we'll see probably more sectors actually adopt this than we actually thought. And one of the neat conversations that we had in the subcommittee for people who are less aware of it, Tim, were, then maybe we should start looking at this to see. So I think, you know, we're not at the stage of hardcore regulation yet, but there's a possibility that much of the regulation will be built around that Risk Management Framework.

MS. HOUSE: Stanley?

MR. GUZIK: So, Nicol and Todd and the subcommittee, I do want to say thank you, you know, a very excellent report. I do think this report comes out, it's actually rather timely. And what I appreciate about the report is about the best practices mentioning existing companies have AI governance frameworks already in place, and this

is, I would say, like an increment or an additional characteristic through its risk governance frameworks.

And then, you know, the way I'm thinking about this document is in the cyberspace, in the software space, you know, financial institutions, we already shifted left, where, you know, the secure software development is, you know, the beginning of the development lifecycle and a design. And what I appreciate about the framework here, especially like, you know, mentioning NIST and the work that NIST is doing, this is an opportunity to think about shifting left with AI governance the way we shifted left with AI security. So it currently falls into those same design patterns that we use to build software.

And then the other part I appreciate about the document is where smaller companies -- we spoke about smaller companies, the challenges that they have with AI. They'll be innovative, developing these new models. You know, they have issues with, you know, access to data and a number of things,

but this also now is documentation for smaller companies to actually shift left and help them build those types of AI governance frameworks. So I would just say very good work. Thank you for this.

DR. LEE: If I could just say something on that. Thank you, because I think that was the subcommittee's point, right, to make sure this is a document that could live in a variety of spaces for the agency and serve as, again, a guidepost for removing I think some of the big burdens that some smaller companies also experience when they don't know what we're talking about with these things. So, yes, we really appreciate that feedback on behalf of the subcommittee.

MR. SMITH: I would just also quickly add, the last point you made in terms of providing some type of information and framework for smaller companies as well, I think that's an excellent point, so thank you.

MS. HOUSE: Yes, I also love that point. The idea of providing frameworks that can help the

smaller businesses and developers address not just security by design and privacy by design, but this idea of like governance and accountability and equity by design, which is all really necessary for building these technologies, especially, I think, with trends towards things like decentralization and other stuff, the ability to meaningfully impact that huge part of the direction of this tech based on the scale of its growth is right now. So that's a really, really poignant point. Thank you so much, Stanley.

Are there any other further questions or comments on this report or other recommendations from either those in the room or online?

Oh, perfect. Nikos.

MR. ANDRIKOGIANNOPOULOS: First of all, I want to congratulate you on the report. I found it extremely interesting for me to read personally. Oftentimes, what we do as Metrika, which is risk management on digital assets, we tend to hear from our customers that digital assets are 24/7, 365. Therefore, when we do risk management, we want to

have as much efficiency, and we want to have as few risk analysts covering the breadth of digital assets and the speed.

So there is a natural use of AI to speed up risk management and allow digital assets to flow in and out of organizations in an easier way. And I think your AI risk framework raises the question on how you responsibly apply AI and you enable the risk management on digital assets.

So I know I'm drawing even more links between digital assets and AI, but I think these two emerging technologies play together in more ways than we imagine, and I think it's great like making those steps or thinking across the board about how there is a hierarchy of risk management that plays into each other. And as finance gets transformed into real time, how do we make use of all of those technologies that really speed it up?

DR. LEE: Thank you, Nikos. We appreciate that as well.

I mean, going back to Carole's point, I think the purpose of the second report is to keep the

build going, right, so that there's a good body of work that comes out of this TAC that actually suggests that these are not technologies in vacuums, and they're not practices in vacuums. They're really interconnected. So I really appreciate that because I think, again, you know, sometimes we do not give credit -- I think someone said this -- to the finance industry for actually being groundbreaking in these areas. And if you hang around at other verticals, we're probably much more advanced, I was just about to say.

I think, Commissioner, we actually have a proposal just to go back to the White House and say we've solved it all because the committee worked collaboratively and did everything I think we've been wanting to do for this country, but it's very nice to hear how it will be very helpful to the industry in weaving all the parts together.

MS. HOUSE: Thank you.

Are there any further questions or comments on the report from those in the room or virtually?

Members, we have now discussed at length the

Emerging and Evolving Technologies Subcommittee report and recommendations regarding responsible AI in financial markets. To further consider these important issues, as there a motion from the body to adopt this report and recommendations and submit it to the Commission?

DR. LEE: You want to do it?

MR. SMITH: Yes, definitely. I make that motion.

MS. HOUSE: Is there a second?

[Multiple seconds.]

MS. HOUSE: So many seconds, thank you.

It has been moved and properly seconded that the Emerging and Evolving Technologies Subcommittee's report and recommendations regarding responsible AI in financial markets be adopted in full by the TAC and submitted to the Commission.

Is there any further discussion?

[No response.]

MS. HOUSE: Are there any further comments from TAC members virtually?

MR. CUTINHO: I have a question, Carole.

MS. HOUSE: Of course, go ahead.

MR. CUTINHO: You know, I didn't know if you were going to do a roll call, but one of the comments I want to make before I vote on this is this. I think the report is a little bit confusing to me personally because it doesn't draw a line or a distinction between artificial intelligence as a discipline and generative AI as a manifestation of that discipline. And if we accept the premise that, you know, this technology is going to be pervasive, then the recommendation that we create yet another governance body or treat this as something distinct from a general technology that, you know, permeates every aspect of business I feel is fraught with a lot of risks and conflicts.

So there are already governance bodies that are risk management frameworks, so I am not very comfortable voting along with the TAC members for acceptance of this report.

DR. LEE: So, Sunil, thank you for that feedback. And, you know, as I've done public service, every report is not for everybody, so we

appreciate and respect you for that.

You know, one of the things that we were challenged with was the extent to which we actually delved into the specific technology cadence, and Francesca had the same conflict that I had as a person who deals on the technical side as well, and the extent to which we also drew the line between how far we were going to go into emerging and evolving technologies in a way that we were going to have very distinct recommendations.

I think for the policy side, at least in my work at Brookings, that we've been real careful about doing that in a variety of reports because the environment is changing. But we would consider it. I think that's the purpose of that first recommendation, that the agency continue to do more work on this, right? So this is not a one-stop deal, that we have a lot more exploration and inquiry to do, but your comment is well noted and something I just wrote down for ourselves to continue to discuss and inquire about and to bring that into the spaces in which we all work.

MS. HOUSE: Thank you. And I'll open this up for any further discussion from the group in response to that.

My own comment, Sunil, I definitely appreciate that, and I certainly have also a lot of feelings about regulation of general-purpose technology. I remember when people were thinking about a blockchain regulator. There's a reason why in the U.S. we took a tech-neutral approach. My reaction to that is mostly that the way that I felt that the recommendation was framed is pointed at adopting a Risk Management Framework, but the same way that NIST issues like guidelines that are unique to particular sectors, that it would make sense for standards are meant to be adapted for use cases.

So I understood it to be something where it's CFTC considering adopting and adapting the Risk Management Framework as would be consistent for its registered entities and for its own use of AI. But I really appreciate that note and would love to open it up for any other further discussion in reaction to this comment. And if not, then we can

move on to the vote.

Oh, yes, go ahead, Todd, and then Jonah.

MR. SMITH: Yes, just one quick comment. We do also make mention of leveraging existing regulations from a risk management program standpoint, so not reinventing the wheel where it's not necessary, but still being thorough and conducting an analysis thereof, so I just want to make that comment.

MR. CRANE: I was actually going to make a similar point about the report itself, and also that, you know, Sunil, I struggle with the same thing a little bit. And as a practitioner in financial services for a number of years, the question has been, so how do we apply current risk management frameworks and governance frameworks when we're using AI? And it turns out it's not that straightforward, and people have really been struggling with it for a long time. And I think I took the report as suggesting some ways to think about fitting within existing frameworks and existing governing structures. And I think the

flexible governance-based approach is consistent with that, so that was my read of it anyway. I think, as a practitioner, I'll find it helpful to connect some dots.

MS. HOUSE: Thank you so much. And thank you, Sunil, as well, for your comment.

Is there any further discussion or commentary from the TAC members?

There's one hand raised, Justin.

MR. SLAUGHTER: Hey, everybody. I'm sorry I had to leave. I'm in the process of getting my daughter, the wages of parenthood.

The thing I wanted to flag -- and I think Sunil's comment is very well taken. In truth, the idea that regulation can ever truly be tech-neutral is something of a fantasy. The only way that works is to assume that all technology is identical when inherently -- and this is very well said by the Europeans -- the nature of technology is to change regulation, vice versa.

That said, I do agree with Jonah that this does strike a middle ground. I think the report

grants that regulation is not tech-neutral because we have to respond and adapt to AI. Finding the dividing line is difficult, but I would remind Sunil this is a report that is nonbinding. It is simply a statement of what we think are best practices and is probably the best way forward to get over that hump of deciding what policymakers should actually do.

So I support this report in part because I think it is not tech-neutral, but it is more open to technology, it is reflective of technology, and that's the better way to deal with it.

MS. HOUSE: Thank you, Justin.

Any other comments from the TAC membership?

[No response.]

MS. HOUSE: The motion on the floor is for the TAC to adopt the Emerging and Evolving Technologies Subcommittee's report and recommendations regarding responsible AI in financial markets and submit the report and recommendations to the Commission for consideration.

As a point of order, a simple majority vote is

necessary for the motion to pass.

I will turn it over to the DFO to conduct a roll call vote.

MR. BIAGIOLI: Thank you, Chair House.

Committee members, when I call your name, I'll start with folks in the room and then move to those participating virtually. Please indicate your agreement with yes, disagreement with no, or you can indicate abstain. Those are the three options. And, as a reminder, abstentions are not counted as a vote.

So we'll start with those in the room. We'll start with our illustrious subcommittee co-chairs and give you the privilege of casting the first vote, so Nicol Turner Lee?

DR. LEE: Yes.

MR. BIAGIOLI: Todd Smith?

MR. SMITH: Yes.

MR. BIAGIOLI: Jonah Crane.

MR. CRANE: Yes.

MR. BIAGIOLI: Nikos Andrikogiannopoulos?

MR. ANDRIKOGIANNOPOULOS: Yes.

MR. BIAGIOLI: Ari Redbord?
MR. REDBORD: Yes.
MR. BIAGIOLI: Carole House?
MS. HOUSE: Yes.
MR. BIAGIOLI: Tim Gallagher?
MR. GALLAGHER: Yes.
MR. BIAGIOLI: Corey Then?
MR. THEN: Yes.
MR. BIAGIOLI: I was off mic there inadvertently. So Corey Then? I heard you.
MR. THEN: Yes.
MR. BIAGIOLI: Todd Conklin?
MR. CONKLIN: Yes.
MR. BIAGIOLI: Cantrell Dumas?
MR. DUMAS: Yes.
MR. BIAGIOLI: Stan Guzik?
MR. GUZIK: Yes.
MR. BIAGIOLI: I don't believe I missed anybody in the room, so we'll move to those participating virtually. Again, I'll go in the order that you all appear on my screen.
So, Sunil Cutinho?

MR. CUTINHO: I abstain.
MR. BIAGIOLI: Justin Slaughter?
MR. SLAUGHTER: Aye.
MR. BIAGIOLI: Christian Catalini?
MR. CATALINI: Aye.
MR. BIAGIOLI: Francesca Rossi?
DR. ROSSI: Yes.
MR. BIAGIOLI: Steve Suppan?
DR. SUPPAN: Yes.
MR. BIAGIOLI: Dan Awrey?
MR. AWREY: Yes.
MR. BIAGIOLI: Michael Wellman?
DR. WELLMAN: Yes.
MR. BIAGIOLI: Adam Zarazinski?
MR. ZARAZINSKI: Yes.
MR. BIAGIOLI: Jennifer Ilkiw?
MS. ILKIW: Yes.
MR. BIAGIOLI: I don't believe I missed anyone. Is there anyone participating virtually whose name I did not call?
[No response.]
MR. BIAGIOLI: Chair House, the ayes have it,

and the motion carries. You have 18 "yes" votes, zero "no" votes, and one abstention.

MS. HOUSE: Thank you so much, Tony. As you heard it, the yeses have it. The motion carries. The Emerging and Evolving Technologies Subcommittee's report and recommendations regarding responsible AI in financial markets has been adopted by the TAC and will be submitted to the Commission for consideration.

It is now time for closing remarks from Commissioner Goldsmith Romero. Congratulations to the subcommittee and the whole team.

COMMISSIONER GOLDSMITH ROMERO: Congratulations, well done. So much work went into this, so much thought. And I've always said constructive collaboration, trusted, respected views that can disagree with each other or take different viewpoints, that is the point of why we do this. It makes for better reports, findings, and recommendations.

And I very much appreciate the weaving, all the AI expertise, the markets expertise, and also

all of the ideas of the interagency cooperation and the rest of the work that the CFTC and this advisory committee needs to do.

So you have just launched a lot of work for yourselves in both reports, and I want to say tremendous progress for the TAC to get out, two very balanced, thorough, thoughtful, foundational reports that are timely and relevant. And the splash is not today. It'll be how the concepts in the report are picked up and looked at with time and as the technology changes.

I will tell you, I know you all read this to vote today, the more you read this report, the more Easter eggs you find that are fantastic thoughts and threads that we can pick up later and do whole reports on.

I'm incredibly proud of this subcommittee. I am incredibly proud of Nicol and Todd for their leadership and for Carole and Ari for their leadership. This was a hard report, as you said, and I hope that we all leveled up today on AI Day. We've heard the same themes. The speakers, all of

their remarks really play into the report, and there's similar themes being weaved all throughout. That's not an accident or coincidence. These are the issues at play in our Federal Government now. I'm always asked is it too early for the Federal Government to do anything?

Since the start of TAC since I took over, we have been looking at AI in a very foundational way so that we can learn and understand the evolution of AI, understand what it's doing in the markets, and understand where it's going. The deep dive into AI governance and best practices, phenomenal. The list of use cases, I know your eyes might run over a table, that is incredible, and just there's so much in here, very thoughtful, very iterative, and very foundational, and I think all of that was reflected in the votes today.

I'm incredibly proud of your work. I'm proud of the trust that you've built together. I'm proud of the closeness that you have to trust each other with raising hard issues and concerns. Thank you, Sunil, for your comment, and it's very much

appreciated, as I always appreciate any diverging views.

So thank you for all of your public service. I told you at the beginning this was public service when I asked you all to serve, and now you know what it means because you've worked very hard on two reports. So it's certainly the best of public service to go into these deep, complex issues and provide meaningful guidance, advice, and recommendations to the Commission. So I'm very proud of this group, and thank you for all of your public service.

MR. BIAGIOLI: To conclude, Commissioner, I have one happy correction to the record. I had announced that there were 18 yeses, zero noes, and one abstention. It's 19 yeses, zero noes, and one abstention.

That concludes our TAC meeting. Thank you, everyone. Meeting is adjourned.

[Whereupon, at 4:36 p.m. EDT, the meeting was adjourned.]