UNITED STATES COMMODITY FUTURES TRADING COMMISSION

TECHNOLOGY ADVISORY COMMITTEE

MEETING

(TAC)

Tuesday,

July 18, 2023

https://www.cftc.gov/PressRoom/Events/opaeventtac071823

Three Lafayette Center

1155 21st Street, N.W.

Washington, D.C. 200581

APPEARANCES [KEY]

CFTC COMMISSIONERS

Commissioner Kristin N. Johnson

Commissioner Christy Goldsmith Romero

Commissioner Summer K. Mersinger [recorded remarks]

Commissioner Caroline D. Pham

TAC MEMBERS

Chair:

Carole House, Terranet Ventures Inc., Executive in Residence

Vice Chair:

Ari Redbord, TRM Labs, Head of Legal and Government Affairs

Hilary Allen, Professor of Law, Associate Dean for Scholarship, American University, Washington College of Law

Nikos Andrikogiannopoulos, Metrika, Founder & Chief Executive Officer

Dan Awrey, Professor of Law, Cornell Law School

Christian Catalini, Lightspark, Co-Founder & Chief

Strategy Officer
Todd Conklin, U.S. Department of the Treasury, Deputy Assistant Secretary of the Treasury for Office of Cybersecurity and Critical Infrastructure Protection
Jonah Crane, Klaros Group, Partner
Sunil Cutinho, CME Group, Chief Information Officer
Cantrell Dumas, Better Markets, Inc., Director, Derivatives Policy
Timothy Gallagher, Nardello & Co., Managing Director, Digital Investigations & Cyber Defense and Chief Security Officer
Michael Greenwald, Amazon Web Services, Global Lead, Digital Assets and Financial Innovation
Dan Guido, Trail of Bits, Co-Founder & Chief Executive Officer
Stanley Guzik, S&P Global Commodity Insights, Chief Technology & Innovation Officer
Jill Gunter, Espresso Systems, Chief Strategy Officer
Ben Milne, Brale Founder, & Chief Executive Officer
Joe Saluzzi, Themis Trading LLC, Co-Founder,

Partner, and Co-Head of Equity Trading
Emin Gün Sirer, Ava Labs, Founder & Chief Executive Officer
Justin Slaughter, Paradigm, Policy Director
Todd Smith, National Futures Association, Director of Centralized Data Science and Analytics
Steve Suppan Institute for Agriculture & Trade Policy Senior Policy Analyst
Corey Then, Circle, Vice President of Global Policy
Nicol Turner Lee, Center for Technology Innovation, The Brookings Institution, Senior Fellow, Governance Studies; Director
Adam Zarazinski, Inca Digital, Chief Executive Officer

Anthony Biagioli, Special Counsel to the Director, Division of Enforcement, CFTC, Designated Federal Officer
Lauren Bennett, Trial Attorney, Division of Enforcement, CFTC, Alternate Designated Officer

A G E N D A

Opening Remarks of Commissioner Goldsmith Romero & Others

Subcommittee Chair Introductions

Responsible Use of AI in Regulated Financial Services

AI Accountability Policy Request for Comment
Presentation:
- Travis Hall, Acting Deputy Associate Administrator, National Telecommunications and Information Administration

Responsible AI
Presentation:
- Nicol Turner Lee, Senior Fellow in Governance Studies and Director of the Center for Technology Innovation, The Brookings Institution

An Impact Assessment of the Proliferation of AI Cybersecurity Capabilities on Financial Security
Presentation:
- Dan Guido, Co-Founder & CEO, Trail of Bits

Break

Regulatory Issues for DeFi, Including DAOs

Enforcement Case Study: Ooki DAO
Presentation:
- Anthony Biagioli, Special Counsel to the Director, Division of Enforcement, CFTC

Extent of Decentralization and Models of Governance in DeFi
Presentation:
- Ben Milne, Founder & CEO, Brale Inc.
- Justin Slaughter, Policy Director, Paradigm

Stability and Security Challenges and Regulatory Implications for Crypto

Presentation:
- Dr. Dan Awrey, Professor of Law, Cornell Law School

Cyber Resilience for Financial Markets

Third-Party Relationships: Interagency Guidance on Risk Management
Presentation:
- Kevin Greenfield, Deputy Comptroller for Operational Risk Policy, Office of the Comptroller of the Currency

Challenges with Understanding Cybersecurity Risk and Implications for Operational Risk Regulation
Presentation:
- Hilary Allen, Professor of Law, American University Washington College of Law
- Timothy Gallagher, Managing Director, Digital Investigations & Cyber Defense, Chief Security Officer, Nardello & Co.

State of Financial Sector Defense and Collaboration to Combat Cyber Threats

Presentation:

- Steven Silberstein, Financial Services Information Sharing and Analysis Center

Closing Remarks and Adjourn

P R O C E E D I N G S

MR. BIAGIOLI: Good morning, everyone. I'm Tony Biagioli. As the TAC designated federal officer, it's my pleasure to call this meeting to order.

Before we begin this morning's discussion, I'd like to turn to Commissioner Christy Goldsmith Romero, the TAC sponsor, for the welcome and opening remarks. Thereafter Commissioners Johnson and Mersinger will also give brief opening remarks. Commissioner Goldsmith Romero?

MS. GOLDSMITH ROMERO: It's nice to see everyone here. I welcome you to the CFTC today.

With artificial intelligence at the forefront of public discussion, significant movements in digital asset markets to decentralized finance (or DeFi) after the collapse of unregistered centralized exchanges and several enforcement actions, as well as escalated and persistent cyber threats, I look forward to the presentations and discussion today from recognized technology experts who serve on the TAC.

I want to thank TAC chair, Carole House, and vice chair Ari Redbord for their leadership of TAC and for putting together today's agenda. I also want to thank CFTC staff Tony Biagioli, who is the designated federal officer, DFO of TAC; our two new assistant DFOs, Drew Rogers and Lauren Bennett; and Scott Lee in my office as well as others in the CFTC who keep TAC running and helped organized today's events.

I'm also excited to recognize our new subcommittee co-chairs. Carole House and Dan Awrey will serve as co-chairs of the subcommittee on digital assets and blockchain technology. Tim Gallagher and Dan Guido will serve as co-chairs of the subcommittee on cybersecurity. Nicol Turner Lee and Todd Smith will serve as co-chairs of the subcommittee on emerging and involving technologies.

I'm grateful for their willingness to lead the work of the subcommittees and all the members who are willing to serve on these subcommittees.

I'm going to first start with responsible

artificial intelligence. Artificial intelligence is at the heart of much public conversation right now including the tremendous opportunities presented as well as some fear of the unknown. Questions swirl around whether we know what we are unleashing.

AI is not new. AI has long been a part of our everyday life as well as a part of financial services. From your Netflix algorithm to your bank's chatbot, Americans come across AI often without thinking about the idea of AI.

At some point, we may have thought the question, am I making this decision because it's something that I want or am I being prompted, but often that may have been a passing thought easily dismissed.

That all changed with generative AI which is new. While there are tremendous benefits of AI, there are growing concerns about harmful outcomes, particularly with generative AI, the list of potential risks to both individual and society may yet be determined -- may not yet be determined.

The concept of responsible AI is also not new.

AI algorithms, logics and -- logic and outcomes should be transparent and explainable in a way that can be audited by humans. Unbiased and representative data has never been more important. Harm should be minimized. It's easy to think of AI enabled market manipulation, fraud and cyberattacks.

There are real concerns of societal harms like bias, abuse and disinformation. And it's important to think through other societal issues such as privacy as well as what types of jobs could generative AI replace and are we losing some aspect of human judgement in those jobs that is important to retain.

The CFTC has an important mission that includes promoting responsible innovation. It is important to increase our understanding of the use of AI in our regulated markets.

When it comes to our regulated entities, we have responsible AI questions and concerns related to organizations' responsibilities. When it comes to AI, including governance, how are decisions

being made and who will make those decisions.

Additionally, there could be greater opportunity for the CFTC to benefit from AI. As a long-time enforcement attorney, surveillance and data analysis immediately come to mind but there could be many others.

I also recognize the benefit of human judgment in these areas, raising the same issues of responsibility when it comes to the Commission's own responsible use of AI.

In continuing TAC's examination of AI, I look forward to the presentation today on responsible AI from TAC member Dr. Nicol Turner Lee, who is a senior fellow in governance studies and the director of the Center for Technology Innovation at the Brookings Institution.

Dr. Turner Lee serves as the co-chair of TAC subcommittee on emerging and evolving technologies along with Todd Smith from the National Futures Association.

TAC will continue to coordinate with others in the Biden Administration evaluating responsible AI.

I look forward to the presentation on the request for comment released in April by the Department of Commerce's National Telecommunications and Information Administration, NTIA, and we welcome Travis Hall, NTIA's acting deputy associate administrator.

NTIA's request for comment advances its efforts to ensure AI systems work as claimed and without causing harm. These efforts build on the blueprint for an AI bill of rights which present -- was presented at the last TAC meeting by Alan Mislove of the White House Office of Science and Technology.

As we consider potential harms of AI, we'll also hear from TAC member Dan Guido, co-founder and CEO of Trail of Bits on the impact of the proliferation of AI cybersecurity capabilities. Mr. Guido serves as a co-chair of TAC's subcommittee on cybersecurity.

Turning now to DAOs and other forms of DeFi. TAC also continues its deep dive examination of regulatory issues related to DeFi.

This examination has been -- become increasingly important as more of the digital asset market is shifting to DeFi and congress is considering additional legislation that includes DeFi. As I said at the start of TAC's examination of DeFi, the central issue is accountability.

We will continue that examination today. Financial regulators are used to central actors. Today we continue the discussion about what decentralization means and the sliding scale that is often DeFi. DeFi is not one size fits all and DeFi can take different forms.

One form is a decentralized autonomous organization, or DAO. Today we'll hear about the CFTC's recent novel enforcement case against Ooki DAO. We'll hear from TAC's very own DFO Tony Biagioli, who was the trial attorney in the case.

We will also have a presentation by TAC members Justin Slaughter of Paradigm and Ben Milne, the founder and CEO of Brale Inc. They will discuss DeFi models, smart contracts and governance.

We will also have a presentation from -- by

professor Dan Awrey on regulating decentralization. Professor Awrey serves as a co-chair of TAC's subcommittee on digital assets and blockchain technology along with TAC chair Carole House.

Finally, cyber resilience. Third party service prider [sic] -- provider vulnerability is one of the top cyber threats. In June, the banking regulators issued their long-awaited guidance on third party service providers. This guidance was lar- -- largely modeled after OCC guidance.

Today we welcome from the OCC Kevin Greenfield, the Deputy Comptroller for Operational Risk Policy, to discuss that guidance.

This is relevant for not only swap dealers who are also sub- -- who many are also subject to bank regulation, but also as best practices for risk management in the financial industry to build cyber resilience and business continuity.

We'll hear from TAC members Professor Hilary Allen and Tim Gallagher on challenges with understanding cybersecurity risks and implications for operational risk regulation. Mr. Gallagher

serves as a co-chair of TAC's cybersecurity subcommittee.

Finally, we'll have a presentation on the state of financial sector defense and collaboration to combat cyber threats by Steven Silberstein of the Financial Services Information Sharing and Analysis Center, the FSISAC.

The FSISAC is a non-profit organization that advances cybersecurity and resilience in the financial system, whose board of directors is composed of cybersecurity executives at financial institutions.

Each of these issues is front and center before us. I really appreciate each of the TAC members' willingness to share your technology, expertise and viewpoints and as always I encourage a broad discussion of a diversity of use today and going forward.

MR. BIAGIOLI: Thank you, Commissioner Goldsmith Romero. We will now have opening remarks from Commissioner Johnson.

MS. JOHNSON: Good morning. Thank you so much

Commissioner Goldsmith Romero, Tony and I think I'm going to get this right this time, Tony Biagio- -- Biagiolio [sic]. Did I get it right? I feel like somewhere you're adding a syllable. Fair enough.

While working on this, it'll be a continued project alongside writing reports and offering of recommendations to this CFTC. It's officially part of the mission of the subcommittee and the commissioners to get Tony's last name correct.

I am grateful for Tony's service as DFO for TAC, which my good friend and fellow commissioner, Commissioner Goldsmith Romero sponsors.

I'm also grateful to each of you here in the room today for the time that you took out of the schedules that are already too demanding to participate and present, to educate the Commission, the staff, market participants and others regarding the need to address these critical issues.

The agenda for TAC today and TAC's subcommittee workstreams address critical issues that will directly influence forthcoming regulation by the Commission as well as future regulation and

legislation.

As was mentioned in the previous TAC meeting, we've discussed and explored President Biden's announcement for a blueprint for an AI bill of rights.

Last month during a speech in San Francisco, President Biden delivered remarks on the risks and opportunities posed by artificial intelligence. The president echoed reflections and concerns that will shape a values driven discourse on the integration of AI in our society.

For those of us who have spent years thinking, researching and writing about the potential and concerns surrounding the integration of artificial intelligence, this is a welcomed approach. In 2008 I began to research and publish literature examining the integration of AI in financial markets.

Not long after, I began to support federal agencies including our Commission and the SEC in the development of regulations concerning or dealing with emerging concerns related to the

adoption of a diversity of technologies that accelerate trading.

Three quick observations: AI technologies tout tremendous promise. A promise to reduce frictions, to effectively enhance efficiency, to permit trading at light speed, yet there are many reasons to carefully interrogate the promises that are made to ensure that we mitigate exposure to potential perils that may also follow from adopting AI.

Let me tell you a little bit about racing at light speed, if I may. Even when technologies offer great promise, we must ensure compliance with existing guardrails including the ability to effectively police AI.

In a paper that I published from a discussion inspired by a symposium panel was su- -- with SEC Commissioner Hester Peirce, I explained federal statutes and regulations regulate risk taking by financial market intermediaries including broker dealers who execute trades and securities exchanges and clearing house platforms where trading occurs.

For almost a century, these statutes have enforced norms that encourage disclosure, transparency and fairness, and modern markets, innovation, and technology challenge these core principles.

The engineering of computer-driven automated trade execution -- I nod my head a bit to Professor Hilary Allen in the corner of the room -- illustrates the necessity of thinking carefully about what it means to automate trade execution, to develop algorithmic trading strategies, to introduce high-frequency trading strategies, and the many accompanying shifts that follow in financial market intermediation.

That academic project trace the evolution and adoption of computer driven trading from the paperwork crisis of 1967 through the stock market crash of 1987, the financial crisis of 2007 and the flash crash in 2010.

Coupled with the accelerated pace of algorithmically driven trading on digital platforms within -- with a global almost instantaneous

internet-based infrastructure, I raised alarms that there's a need for guardrails that effectively address any integration of high frequency trading, permitting co-location and particularly with respect to our ability to effectively resolve enforcement questions.

In the paper I focused on one specific case study that illustrates two strategies that really ought to be the focus of a lot of our conversations: front running and spoofing.

With more time I described to you that I focus in the paper on the flash crash and the SEC and CFTC enforcement divisions' initial conclusions regarding automated algorithms and the extent that they played a role in the flash crash.

I also explore but -- that by 2015 the Department of Justice and CFTC investigations revealed that a rogue London-based trader, we'll call him the vendor for now, had manipulated the E-mini S&P 500 by using an algorithm to flood the CME with sell orders for E-mini S&P 500 stocks. I'd say lots about spoofing, I'd say lots about front

running with more time but I'm sure that will come as well.

I'd like to wrap up with just a few reflections on understanding the ethical implications of AI. I applaud Commissioner Goldsmith Romero and the TAC subcommittee focusing on the implications of AI.

As an associate dean and a DAO professor at Tulane University, I convene computer scientists, programmers, engineers, market participants, lawyers, public interests advocates and academics for a day long symposium examining the implication of AI NHS Society.

From that discussion I published a paper in the Journal of International Comparative Law where my co-author and I examine a diversity of areas in our markets where we're integrating AI and the need to be thoughtful that this awesome barrage of technologies that we are integrating into our society really are fit for purpose and are consistent with the rules we've long adopted.

Finally, I supported the Administrative

Conference of the United States in a project that I hope will also ultimately influence regulation here at the CFTC. ACUS developed a number of reports following a consultation. I was one of a number of consultants that supported ACUS examining the potential and perils of administrative agencies subject to constitutional protections integrating AI.

As reported in our final agency report here at the Commission, we are already integrating AI in a number of ways and anticipate further integration in due course. I'm deeply thoughtful about the extent that we integrate AI.

I'm even more deeply thoughtful about the reality that even our efforts today may not be sufficiently comprehensive. I'm mindful that the integration of AI has implications not just for large financial institutions but there are many implications for individual market participants or consumers as well.

Here questions of bias and the potential for privacy concerns really take root. Just three years

ago -- oh no, it's four now, four years ago a week from today, I testified before Congress specifically describing concerns around the integration of AI in the context of residential mortgage markets. While AI might present as neutral in many contexts we must be sure that the outcomes are free from bias and fair for everyone.

I'm also excited to hear from the other subcommittees that are presenting today and due course will have lots to share and engage with you regarding digital assets and the oversize digital assets in our markets and cyber resilience as well.

Thank you all so much for coming. Thank you for having me, Commissioner Goldsmith Romero.

MR. BIAGIOLI: Thank you Commissioner Johnson. Commis- -- Commissioner Mersinger has prepared recorded remarks.

RECORDING OF MS. MERSINGER: Good afternoon and thank you to Commissioner Goldsmith Romero for hosting another meeting of the Technology Advisory Committee. I regret not being able to join today's meeting in person but I have no doubt that the

presentations and discussions during today's meeting will be educational not [inaudible] but to everyone on the committee, those participating in today's presentation and anyone who's joining us whether in person or virtually.

As some of you may notice -- have noticed, July has become the month of advisory committee meetings here at the Commodities Fuder [sic] -- Futures Trading Commission. With four separate committees holding meetings almost back to back and actually five if you count the Energy and Environmental Markets Advisory Committee meeting that I sponsored at the end of June.

Now, as a regulator who has to remain up to speed on a wide variety of topics, all at the same time, these advisory committee meetings are a great opportunity to hear directly from experts and learn about matters important to the work we do here at the CFTC.

The information and discussions occurring at these meetings go far behind anything we could obtain on our own, whether it's to reading reports,

write papers, research, et cetera. But as a Commissioner who also enjoys sharing a few remarks at the start of these meetings, I have to admit I'm running out of material.

As this statement is before the advisory committee devoted to technology, I decided to enlist the help of a few people in my life who are very adept at using technology: my children. Not only did they provide awesome ideas for my statement, they also confirmed that this was the first time anything about my job has been interesting to them.

So I asked three of my four children for their thoughts on artificial intelligence in decentralized findings. Now, my panel was missing one child because my 16-year-old daughter now has her driver's license, she is never home and cannot be bothered by conversations with family members.

I also want to just give you a quick warning that the expert panel included two elementary school age boys. So there were a lot of comments that I deemed that were not appropriate for public

consumption. So I will share some of the comments that did make the cut.

When asked about AI, my youngest son told me he has no idea what that means. And so I thought I'd offer a simple explanation. I said, it's the computer thinking on its own, making conclusions based on information it has available. Now, my son is an avid football fan and his response to my inquiry and simplistic explanation of AI was a pretty shrewd test question of his own.

He said, so when I ask the computer who is the best quarterback in the NFL, and it says Patrick Mahomes, is that because the computer thinks it should say Patrick Mahomes is the best or because the computer knows that he is my favorite quarterback? Or is this just a fact that Patrick Mahomes is the best quarterback?

Now, my 12-year-old son and 14-year-old daughter were a little more familiar with AI. In fact I learned a new term from my 12-year-old son -- non-playable characters, or NPCs, as they are known the gaming world.

I asked him to explain NPCs to me and how they use AI. His response was, mom, you wouldn't understand. All right, fair enough. Now, my 14-year-old daughter offered another example of AI that was also new to me, this time in the context of social media.

She told me that she likes to text SnapChat AI because, according to her, it always reads her text, it texts back immediately saying something that she would want to hear and it will never ghost her list her last boyfriend did on SnapChat.

Now, the conversation regarding DeFi were much less substantive. DeFi seems to be a concept that they have yet to fully grasp.

My youngest son when asked that he doesn't know what that means and quickly ended the conversation by returning to his video games. My 12-year-old son said, you mean like crypto currency and those gorilla NTFs [sic]? NFTs? That's dumb. Why does anyone use real money to buy fake money?

I followed that up with an example suggesting that when he buys tokens on his favorite video game

with his birthday money he is essentially buying fake money with real money. So with that he responded, yeah, but that's different because I can buy cool weapons.

And my 14-year-old, said, what, do you mean like Venmo? I like Venmo because you can like send me money whenever I want it. Now, I do not report these admittedly very unscientific survey results to be clear about the issues that are on the agenda for today.

That really it's just the opposite. AI and DeFi, and I'll add cybersecurity to that as well, are not hyperical- -- hypothetical challenges of futures. They are challenges we are facing now, especially in our financial markets and the -- and challenges for those who regulate them.

Yet they often come wrapped in new concepts and terminologies that seem just a thorn to many of us as they are to my young kids. And thankfully today we are going to hear from real experts who will discuss, explain and analyze these important topics and more.

I am certain we will all walk away from today's meeting much smarter and much better prepared to consider the impact of these technological advances on the markets we regulate at the CFTC.

I'm also looking forward to reviewing any work product from the three TAC subcommittees: The digital assets and blockchain technology subcommittee, the emerging and evolving technologies and the cybersecurity subcommittee, especially in areas where we may need to take a second look at our regulations to make sure they remain fit for purpose.

As always, thank you to all of the CFTC staff who make sure these meetings are a success. I believe I speak for my fellow Commissioners when I tell you that these meetings cannot happen without you.

Again, thanks everyone for being here today and I'm looking forward to very interesting discussions.

MR. BIAGIOLI: Thank you, Commis- --

Commissioner Mersinger. Commissioner Pham is live from somewhere in the world, not here, and so she will deliver her remarks now virtually.

MS. PHAM: Thank you. Thank you so much for having me. I'm -- it's a pleasure to be able to speak to you all today.

I thank Commissioner Christy Goldsmith Romero and I'm pleased to support her sponsorship of the CFTC's Technology Advisory Committee. I would also like to thank the TAC Designated Federal Officer Anthony Biagioli, and Alternate Designated Federal Officer Lauren Bennett, and other CFTC staff for their work preparing today's meeting.

I welcome each of the TAC members as you explore timely issues regarding responsible artificial intelligence, the centralized finance or DeFi and cyber resilience. Thank you Commissioner Goldsmith Romero for your leadership. I also congratulate the Division of Enforcement for their work on the Ooki DAO case.

I wish that I could be with you in Washington but I am still in New York City following my GMAC

meeting yesterday as I know Commissioner Goldsmith Romero was in D.C. yesterday and dialed in for my GMAC meeting.

Recently I specifically addressed AI and other technological advancements that may impact financial markets. Across the industry, risk professionals have a critical role in safeguarding our markets. I discuss the importance of utilizing existing risk governance frameworks and risk management discipline to identify, measure, monitor and control emerging risks and new technology.

For example, operational risk management includes technology risk, cyber risk and third-party risk. Model risk management is key for AI risk governance. Businesses must also consider strategic risk and compliance risk in light of recent technological development.

I recently stated that our registrants must be vigilant and address new and emerging risks through various risks stripes as appropriate; whether from changing market conditions, technological developments, geopolitical concerns,

or any other event.

At the last TAC meeting, I remarked on the many years of work by policymakers such as the Financial Stability Board, the Basel Committee on Banking Supervision, the International Organization of Securities Commissions and other regulatory authorities around the world to implement laws, regulations, and standards for operational resilience.

Regulated entities including the vast majority of our swap dealers and FCMs that are banking organizations have implemented comprehensive enterprise wide operational resilience programs.

Operational resilience, as noted by U.S. prudential regulators in 2020, encompass governance, operational risk management, business continuity management, third-party risk management, scenario analysis, secure and resilient information system management, surveillance and reporting and cyber risk management.

As you can see, cyber risk or cyber resilience

is only one component of an operational resilience program.

It is my view that the CFTC's approach to cyber risk or third-party risk should appropriately recognize that these risks are within the discipline of operational risks and that all of these risks are part of, but not the same as, an operational resilience program.

I look forward to hearing from Kevin Greenfield, Deputy Comptroller for Operational Risk Policy at the Office of the Comptroller of the Currency on the recent interagency guidance on third party risk management. Many of our swap dealers are OCC chartered national banks and it is essential that the CFTC understands the prudential regulation of banking organizations.

I appreciate Commissioner Goldsmith Romero's engagement with our fellow U.S. regulators on these issues and for her leadership. The insights and perspectives shared through the TAC's work will help to shape the CFTC's approach to new and emerging technologies.

My thanks again to Commissioner Goldsmith Romero, the TAC members and speakers for your time and commitment to fostering responsible innovation in our market. Thank you.

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 your microphone is on when you speak. The meeting is being simultaneously webcast so it's important that your microphone be on so that the webcast audience can hear you.

If you would like to be recognized during the discussion, please change the position of your placecard so that it sits vertically on the table or raise your hand and the chair or vice-chair will recognize you and give you the floor. Please note that the webcast cannot be muted. I repeat, the webcast cannot be muted, so at breaks, please turn off your microphone or else it'll pick up any -- potentially pick up any side conversation.

If you're participating virtually and would like to be recognized during the discussion for a question or comment or need technical assistance, just message me within the Zoom chat. I will alert Chair House or Vice-Chair Redbord that you'd like to speak during the discussion period that follows the prepared remarks and presentations.

Please identify yourself before speaking, just say your name for the benefit of the court reporter, please. Please also speak directly into your microphone for optimal audio quality on the webcast.

If you're -- please make sure you unmute your Zoom video before you speak and mute both after you speak. Please only turn on your camera when you're engaging in discussion 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, I'd like to do a quick roll call of the members participating virtually so we have your attendance on the record. After I say

your name, please indicate that you are present and then mute your line. Christian Catalini?

MR. CATALINIE: I'm present.

MR. BIAGIOLI: Jennifer Ilkiw?

MS. ILKIW: [no audio response].

MR. BIAGIOLI: Steve Suppan?

MR. SUPPAN: Present.

MR. BIAGIOLI: Michael Greenwald?

MR. GREENWALD: Present.

MR. BIAGIOLI: Gün Sirer?

MR. SIRER: Present.

MR. BIAGIOLI: Dan Awrey?

MR. AWREY: Present.

MR. BIAGIOLI: Michael Shaulov?

MR. SHAULOV: [no audio response]

MR. BIAGIOLI: Francesca Rossi?

MS. ROSSI: [no audio response]

MR. BIAGIOLI: John Palmer?

MR. PALMER: [no audio response].

MR. BIAGIOLI: And Joe Saluzzi.

MR. SALUZZI: Present.

MR. BIAGIOLI: At this time it's my pleasure

to turn things over to the TAC -- to the TAC chair, Carole House.

MS. HOUSE: Thank you so much, Tony. It's my pleasure to introduce again, the newly appointed co-chairs of the TAC subcommittees and thanks to all the members of the subcommittees as well. The list of membership has been published on the CFTC website. Dan Awrey and I will serve as the co-chairs of the subcommittee on digital assets and blockchain technology.

Tim Gallagher and Dan Guido will serve as the co-chair of subcommittee on cybersecurity. And Nicol Turner Lee and Todd Smith will serve as the co-chair of the TAC subcommittee on emerging and evolving technologies. So we'll be turning it over to all the co-chairs to give some brief introductory remarks. We'll start with Dan and myself.

So first I'll -- I'll say that it's -- it's an honor after having the opportunity help drive President Biden's efforts in ensuring responsible development in digital assets to now get to serve

on -- and not just as the chair of the TAC but also as co-chair with someone who I respect so deeply, Dan Awrey is my other co-chair and with an incredible expertise across government academia, TradFi, DeFi, fintech, all represented on the subcommittee for digital assets. This is a wonderful way to serve and I'm honored to have been selected and asked by the Commission.

I'm sure all of noticed that a lot of the issues that Commissioner Goldsmith Romero spoke to when she was talking about AI issues around bias and security and abuse and privacy, those sound very familiar to us and everyone dealing in the blockchain and DeFi space and ultimately it points to the issues that I know all of us share and we talked about in our first meeting, that responsible innovation does not mean unchecked technological advancement without regard to societies and democracies and consumers and businesses.

Ultimately this goes back to many key issues but especially accountability, which I know is emphasized by each of the Commissioners in their

remarks and something that I'm sure will be the focus of much debate and fun discussion during the subcommittee meetings.

But these issues will require vigor and rigor and intellectual honesty and timeliness as well -- as well as open minds to examine the reality of these technologies as well as the great potential for promise as well as peril without the right protections being put in place.

So I look forward to the hard work that is coming with the subcommittee. Looking forward to working with most members of the committee, all members of the subcommittee on digital assets, so -- and blockchain technology, it's going to be a real -- some really exciting work.

So Dan, I'm going to turn it over to you for introductory remarks, my fellow co-chair.

MR. AWREY: Thank you so much. I will keep it very brief because I think Carole has framed the work of the subcommittee brilliantly.

I'm really excited to work with Carole and the rest of the subcommittee and I'm especially excited

because of the opportunities that we're going to have to identify and debate some important foundational questions and both the regulation of Defi. I -- and then to use that as a building block I think for understanding and evaluating potential past forward for regulation.

So thank you again and I look forward to engaging with all of you on the subcommittee's work.

MS. HOUSE: Thank you so much, Dan. And Tim and Dan, will have some introductory remarks regarding the subcommittee on cybersecurity.

MR. GUIDO: Thanks. Hey. I am Dan Guido, the CEO and founder of Trail of Bits. Very excited to be here and to help the CFTC and the financial industry understand what new opportunities look like for defense as well as make them aware of upcoming risks that may impact their cybersecurity.

Very excited to be working with Tim and help translate these into efforts at good, well-crafted policy. So yeah. Thank you for inviting me and happy to be here.

MR. GALLAGHER: Hello. Tim Gallagher here, managing director and chief security officer and CISO at Nardello & Company.

Thank you, Commissioner Goldsmith Romero, for the opportunity to be on this subcommittee where we have the opportunity to surface some critical threats to the system and ways that we can try and mitigate them, push together -- push forward some ideas and start to socialize them and hopefully help move the ball forward in protecting our infrastructure.

MS. HOUSE: Thank you so much, Dan and Tim. Really looking forward to that work. And now finally, Nicol and Todd, some introduct- -- introductory remarks regarding the emerging and evolving technology subcommittee.

MS. TURNER LEE: Thank you very much. I'm Dr. Nicol Turner Lee, Senior Fellow in Governance Studies and the Director of the Center for Technology Innovation at the Brookings Institution. I'm excited to be here.

Thank you, Commissioner, for allowing me to

participate in this esteemed committee, primarily because this is groundbreaking. This is frontier work, and on the Emerging Technologies Committee, my colleague Todd and I intend to try and stay ahead of the curve of groundbreaking, which is often difficult to do because technology always moves.

So we're looking forward to advising on some of the sociotechnical concerns as well as those interests that are related to the public and how we can better serve and protect them through those technologies. So I'll pass it over to Todd.

MR. SMITH: Thank you, Nicol. My name is Todd Smith. I'm Director of Data Science and Analytics at NFA, I also chair NFA's Data Governance Committee, and lastly, I lead NFA's innovation efforts.

I'm very thankful and appreciative of the opportunity to serve on TAC with you all and especially appreciative to be co-chair with Nicol on the Emerging and Evolving Technologies Subcommittee. So thank you very much.

MS. HOUSE: Thank you all so much for those great remarks. And I love that, frontier work. That's very exciting. Great way to frame it, Nicol, and the expertise in our co-chairs as well as all the members of the subcommittee is going to bring some incredible -- some incredible focus and interesting perspectives into the work of the subcommittees. Thank you.

So moving onto our first topic of the day, which will focus on the responsible use of AI in regulated financial services.

It is my pleasure to introduce our first presenter regarding AI issues, Travis Hall, who is Acting Deputy Associate Administrator of the National Telecommunications and Information Administration, NTIA, at the U.S. Department of Commerce, who will now present NTIA's April 2023 request for comment on AI system accountability measures and policies. Travis, over to you.

MR. HALL: Hi. Thank you so much for having me. Really appreciate the presentation so far. For those who aren't familiar with NTIA, we are a small

agency within the Department of Commerce. Our -- you might -- those who do know us tend to know us for broadband. We're responsible for the Internet for All broadband grants. We're also in charge of managing federal spectrum.

My office, however, is part of our policy shop, which is actually divided into two, the Office of Policy Analysis and Development and the Office of International Affairs, where we are, in terms of our statutory mandate, the president's principal advisor for telecommunications and information policy.

And so in that role, we have been kind of a little bit of a think tank within the federal government. The history of that statute is that we actually used to be within the White House and then got taken out and smashed together with spectrum back in the '70s, which is why we have that kind of like statutory direct link.

But we're a bit of think tank thinking through some of these harder issues on telecommunications and information policy starting in the '80s. But

more recently, folks probably know us from our work on consumer privacy, such as the 2012 Consumer Privacy Bill of Rights and then the draft pledge that came out in 2015 and our work on 5G and other types of reports.

But in that vein, we are taking a look at artificial intelligence and specifically we're looking at accountability policy. And I know that accountability came up quite a bit here. And what we're looking at, because there's a lot of conversations around to what AI needs to be held accountable; right?

What does it mean to be trustworthy? What does it mean in NIST parlance; right? What does it mean for systems to be unbiased, to not cause harm in particular circumstances?

But what we're focusing on is we're actually taking a little bit of a step back and saying, how do these entities get held accountable? Somewhat agnostic to what you are holding them accountable for.

Do we actually have the mechanisms in place in

terms of auditing ecosystems, impact assessments, transparency that -- what the -- what Professor Ellen Goodman, who's on detail with us leading this work, calls the plumbing of accountability; right?

So that is what we are focused on in this. We are of course thinking through in terms of our work within the interagency, in terms of some of our broader conversations, the questions around that like, to what; right?

Like, what -- how -- what does trustworthiness mean? How -- like what are the standards? Things like that. But in this particular effort, we're really focused on the how. And so as you can see on the slide, right, we're looking at the tools, the definitions, the inputs, the resources. What do we need?

And we can go ahead and go to that second slide. So and for the most part, we are focused a bit on the trustworthiness as kind of the benchmark of what we're looking for. Things that entail choices about design and documentation, risk allocation -- we already had conversations about

that -- and of ultimately regulation and enforcement.

However, I will say that these tools are not just for regulators. They are also for entities, companies, purchasing services from other companies; right? In order to understand what you're buying is not just simply somebody sitting on the other side of the computer with a magic eight-ball; right?

That they're act- -- that you are actually getting the services and the things that you are requesting and requiring. That goes for private companies. It goes for individuals procuring services. It goes for the government procuring services as well.

And ultimately, since we are a policy shop, right, we're not a technical standards development body, we're not a regulator; we aren't focused on what are the policies that either within the federal government or in terms of statutory reform or rules that regulators could -- regulatory principles that regulators could take on or think

through, what the policies are, in order to foster an ecosystem.

Because I do think that at this stage, artificial intelligence isn't new; right? Let's just put that out there. Artificial intelligence, we're talking about some of the frontier models, some of the large language models, things like that, that are pushing boundaries and doing new things.

But we've been talking about artificial intelligence for almost going on a decade. And before that, we were talking about big data; right? And before that, we were talking about other types of issues, algorithmic issues, algorithmic justice, things like that.

So these are ongoing conversations with permutations that with these new frontier models do complicate things further; right? The frontier models do take some of the arguments that we were having or trying to think through some of the tough nuts that we were trying to crack in terms of black boxes of these algorithms that makes them even

blacker; right?

And so that I feel like that is something to put out there and say that we are trying to think through how to do this, but we're not building completely from scratch here. And we could go to the next slide. So what we do when we do our work streams is we rely heavily on public input.

So we put out a request for comment with 34 plus questions on a range of topics. I'm not going to read the slide to you. Or conversely, the request for comment itself. But asking really what are the objectives that need to be put in place? What are the existing resources? And, ultimately, what is needed in terms of inputs, in terms of, what are the barriers?

And at the end of it, like, the real gems and the comments that we're finding are, what are the policies that we -- that we should be putting in place? Should these things be mandatory?

Should they be voluntary? Should they be hor-- you know, across the board, kind of like the EU AI Act is looking at artificial intelligence as a

single thing? Or should it be sectoral in approach?

The EU AI Act is taking a risk-based approach but it's still overall a broad scale AI regulation. And what are some of the harder questions? What are some of the harder things that we do need policies on? Like researcher access; right? Or auditor access. Like, is it -- should we be reliant on internal auditors or external auditors? And what are the legal protections that are involved?

So we're very fortuitous in our timing. We he-- our request for comment was in process for I would want to say maybe six months, seven months. But it came out right at the time that GPT-3 was released, and so we ended up getting over 1,400 comments in our process.

So we're in the process of reviewing those comments in order to hopefully have our report out sometime in the fall. A little bit of a preview of some of the -- some of the insights that we have gotten from some of the major commenters.

One, there is a focus on application deployment and risk. And that, I think, is a -- is

in general an area of consensus, that broad scale flat regulation of all AI doesn't really make sense.

And that comes both from industry and from folks who are academics in civil society that there really should be a focus on specific deployment and use. However, in addition to focusing on deployment, the auditing and accountability should -- there should be a look of view towards the entire life cycle of an AI development; right?

And we did -- there was a bit of an emergence of, in terms of accountability tools, that you can have accountability tools for the models themselves, but you can also have accountability tools for the governance structures.

And that in many instances, in terms of the development of the AI, where you have rapidly developing models, that actually the governance structure is what you should be auditing towards or doing impact assessments of, rather than the individual models, to ensure that the companies have the proper governance structures in place.

And then I will say that there were also some areas of stark disagreement, one which you -- probably won't surprise anyone is about whether audits should be mandatory or not.

There was also quite a bit of discussion and hard discussion of questions around liability, particularly in terms of access, in terms of trade secrets, and in terms of other areas where, in order to hold these systems accountable, particularly if you're looking at third-party auditors, those auditors do need to have a certain degree of access to the crown jewels, such as it was, of these systems.

And that that can be very -- that companies are very resistant to that. They're very concerned about that and they're very concerned about potentially not just simply, you know, handing it over to the government is one thing, but handing it over to a third-party accounting firm or something like that seems -- is a bit more concerning for them.

I will say that there was some callout and

callback to other forms of auditing impact assessments, etc., conversations around privacy impact assessments, SORNs, things like that, as well as of course traditional financial auditing as models to look towards, but certainly not necessarily one to one replicate. So those are just some of the nuggets that I can talk ab- -- that I can -- that we can pull from the comments. Certainly our comment summary is not complete.

We're still working through the vast majority of the comments, but we are looking to have our report in draft form sometime by the end of the summer and hopefully publish sometime this fall. And with that, I'm more than happy to answer any questions.

MS. HOUSE: Thank you so much, Travis. We're going to go straight into our second presentation then have a great discussion time to bring both you and Nicol with some questions.

So for our second presentation regarding AI issues, Nicol Turner Lee, Senior Fellow in Governance Studies and Director of the Center for

Technology Innovation at The Brookings Institute, will present on the topic of responsible AI. Nicol, over to you.

MS. TURNER LEE: Thank you. Thank you again for having me and, again, thank you, Commissioner, for leading this really extraordinary effort and having me be a part of it. So in my role at Brookings, I pay attention to a variety of issues, one of them being artificial intelligence, in particular bias mitigation, with a interest in how we make it more inclusive and equitable and responsible.

And so I'm privileged to actually get this presentation because this is something that, over the years, as it was mentioned earlier, we've seen this outgrowth of interest in AI. And I think it's important for us to do some level setting around that.

So thank you to Travis for the work that NTIA is doing alongside many other federal agencies that are providing guidance. So I think there's a value in the work that we're doing today. Next slide.

I'd like to start with this slide to give an idea of this new ecology, which is interesting, because for those of us who are familiar with this topic, artificial intelligence, the big bubble, is all about autonomous systems, whether they be autonomous vehicles, decision making assets that we may have that are empowered by machine learning algorithms, which were mentioned earlier by Commissioner Johnson.

When we think about the power of repetition, the predictability of machines are just much greater than what we've ever imagined since IBM's Watson machine.

And if we add in deep learning, which applies deep neural networks, we take this autonomous environment in which we're in and we bring in these other signals, eye contact, voice, image, that have implications when it comes to hiring capacity, suitability, etc., elections, political misinformation being one fact.

But it's again, as you see, this evolution of this ecology. If you notice, I put generative AI

chat boxes in a box because it is still new. When I did this presentation a couple years ago, I couldn't put it in the circle because it was not spoken about. But the idea is chat boxes are joining this new ecology and it's something taken together.

We have to think about the implications of these emerging technologies. I would like to say, so not to appear to be a pessimist, that these technologies are not just for us to interrogate them for concerns.

There are obviously great efficiencies that come of these technologies. The COVID pandemic was one of those. We actually leveraged artificial intelligence for vaccination development. It allowed us for emergency authorization. My point is it's like a two-sided coin. There will be opportunities on one side and perils on the other.

For the purpose of my presentation, next slide, I want to talk a little bit more about bias because I think this is why it's actually taking the domain and interest of public policymakers. For

the most part, as it was mentioned by Travis, we're talking about the life cycle of algorithms or autonomous systems.

We're really no longer talking about when it's developed and designed. We're talking about the iterative process in which these models are actually deployed. And we can actually apply this model from general autonomous systems like vehicles all the way down to generative AI models.

What do I mean by that? Who designs it is really important. What that research question is, who is sitting at the table when that question is designed, how do we implicate the concerns before the model is actually deployed out into the real-world matters.

And in a lot of my work, there is not a lot of diversity, diversity of people of color, women, sociologists like myself who sit at the table at these more generally purposed technologies. In addition to that, training data. And this is really important for this group.

I know some of you may have heard this but you

haven't heard my presentation on this. When we think about statistical differences and human prejudices that are embedded in training data, it's something that we should take concern over.

Dr. Renee Cummings at the University of Virginia calls this "traumatized data" -- data that we actually scale and mine from publicly available sources that come with their own in-baked prejudice -- prejudices.

And that's important for us to consider because in some cases, they're overrepresented, in some cases they're underrepresented.

It was mentioned I think in some of the opening remarks about criminal justice algorithms that stand out for many because the training data is often based on people who are part of the criminal justice system, part of mugshot databases, part of some type of identifiable exchange with an institutional authority.

And when you begin to look at how we train algorithms based on determining someone's sentencing or relief when it comes to bail, the

training data tends to skew disparately for people of color. And that is again a very important aspect for us to consider when we look at these models. Because in the space that we're going to talk about these models, where does it skew?

And that's always a question that stays in my mind. What are we looking at when we think about the data that these models are being trained on? And that very much impacts results.

There are always going to be trade-offs in computational models. One, who's at the table determines the inclusiveness of that model, the data in which we're training them, and how we interpret those results.

So interestingly enough, if I go back to the criminal justice algorithm, there's a high likelihood that judges will be right and there'll be a certain level of validation. But that likelihood sits on top of a fractured and disparate criminal justice system that already overincarcerates and incriminates people for false arrest.

My point is, when we look at these models, as policymakers, we really need to be concerned about this stage. And just in response to my colleague here, Travis, in terms of the black box, for me, the black box is clearly apparent. Because we know those things that we're constantly navigating around based on the historical trauma of these contexts in which they're deployed.

One other thing I'll just note is -- put in your attention is some work by Michael Kearns around the ethical algorithm where he actually identifies inferences, whereas people do not necessarily consent to these models actually using their data.

But it's inferred because, as I've heard from my technologist friends, what started as 10 variables or attributes has now become millions of attributes around individual people from your biometric information, to the time that you get online, to what my colleague Aaron Klein says at Brookings, the type of device you use determines your credit worthiness.

My point is these models are not in a vacuum in isolation. They're part of an extension of the society in which we live. Next slide.

And they show up in a variety of places, whether it's ad targeting, employment biases I've mentioned, facial detection errors and inaccuracies. I also sit on the National Academies of Science Commission that was put in place by President Biden to look at the use of FRT, facial recognition technology, in law enforcement.

We've seen improvements but we still see false arrests or false use. Search query misrepresentation, predictive policing as I've mentioned, credit scoring and other financial service errors, political disinformation and misinformation, and most recently we're starting to see more of this in healthcare practices and research, contributing to further health disparities. Next slide.

I wanted to share this before I go into our particular concerns in terms of the regulatory landscape for responsible AI. Obviously, to get to

a place of responsible AI, it requires a sociotechnical approach.

As a sociologist, of which I tell most people, if you don't have a friend like me who is a sociologist, you're missing out. We make great dinner conversation partners. We really have to look at biased decision-making. And this is not to suggest, my friends, that there will not be moments of differential treatment of people.

That is why the technology works so well, because it knows us. It knows the recommendations of our movies, as our Commissioner said. It knows where we want to vacation. But it also can use that data for biased decision making, emboldening gender, racial and other implicit biases that reinforce systemic discrimination.

It can be used by malicious actors who create deepfakes, conduct unlawful and unauthorized surveillance or profiling. And more, we've spoken about it, hopefully we'll speak about it more, job displacement, data privacy, copyright issues, carbon emissions and environmental impact due to

these large training models. That's a societal side of it.

And technically how we operationalize these and measure these values of fairness continue to be elusive, especially when we talk about ethics. And how we envision the technologists in charge to actually make those trade-offs are equally important, as well as the standards that we set.

And now my new thing, this balancing and conflation of AI efficiency with disparate outcomes. We want AI to be productive. It's helping lawyers organize large case laws and studies. It's helping researchers mine different large datasets for quantitative research. But at the same time, we're still seeing disparate outcomes. Next slide.

I'll just share this because I think that there's a lot to be gleaned in terms of our learning context, and I'll go over this quickly, which is just combining human values with artificial intelligence so that we do not have them in the abstract, whereas you see things like rights, ethics, law, privacy, fairness are really

key to this on the other side of statistical measure. Next slide.

And here, for the purposes of this committee, my hope is that we'll look at this around a variety of use cases. I'm new to this particular part of the financial sector, but I know risk management is pretty key when we look at compliance protocols, credit lending decisions. Fraud prevention is equally important.

And there's been a history already of algorithmic trading which has had some success and has come with some lessons when we think about how we use these models, again, that have been trained on particular datasets with interpretive results, how we use those to maintain lawfulness.

And obviously generative AI, as we engage it for -- with clients, investors, research, as well as drafting contracts reports and presentations. Next slide. And the key challenges overall in financial services and products, more generally.

I just shared with the Commissioner's team, with Anthony, a report that I came across from the

GAO who talked about fintech. But most importantly, it's getting at this bias in and bias out problem, keeping more humans in the loop, but most importantly, transparency.

And what's interesting, and I'll go into this as I wrap up, is the transparency around the explainability of these models to people who are non-technical. But more importantly, the disclosure of their application to everyday citizens who, by the way, still do not have here in the United States the relevant and necessary privacy protections to know what these models actually do.

And there's always the avoidance of the type of risk that comes with that. So like I said, two-sided coin where there are opportunities and perils. I'd like to just close my remarks with providing some landscaping of the regulatory environment which I think is really important to this conversation.

Obviously right now, we're seeing a couple things happening. So when Travis and I use this word frontier, it's actually a new thing with

regulation that just came out a few weeks ago that people are talking more and more about. Took an academic to sort of posit that.

But I like to consider it a combination of soft law, where we're seeing a lot more movement towards voluntary commitments, picking up on some of the great work of coalitions as well as what we've seen with the White House. We're seeing voluntary self-regulation being more so le- -- more so towards promoting innovation with minimal disruption to business models.

So we've seen this in technology. I've been in technology for 30 years. We're constantly balancing innovation regulation. But on the voluntary commitment side of soft law, I think we're seeing more collaboration and consensus, primarily led by industry. Obviously, the limitations with voluntary commitments are no enforcement mechanisms or legal remedies when there are violations.

And then on the other side of the continuum is hard law where we're seeing more enforceable requirements and regulations. We're seeing that

outside of the United States, and I'll talk about that just briefly, and we're seeing this push, as you all have -- may have noticed, which I find to be interesting.

President Biden and The White House recently came out with more soft law voluntary commitments. We're seeing now Senator Schumer talk more about hard law, which is enforceable regulation here in the United States which I'll talk about. But that requires bipartisan support as well as the right oversight at the state and federal level. Next slide.

I thought I would throw this up on the slide just to show you how our counterparts in the EU are dealing with regulation more generally, whether it's the general data protection regulation or my colleague Alex Engler at Brookings speaks about different approaches and different environments.

We're seeing tiered systems. The recently proposed AI Act, which actually just got passed, I believe it is in the -- past the comment period, is actually looking at mandatory disclosures. This

whole idea of labeling is very much in that regulation.

High-risk models are actually on the top of the list for our EU counterparts when it comes to providing the standards, technical documentations and really banning things that look like they have unacceptable risk. Very prescriptive as we all know. When we remember when the GDPR came out, that was six years ago.

We still don't have federal privacy legislation here so that tells you that we're still trying to work through a more multi-stakeholder collaborative process. And then obviously the EU has taken the lead through the Digital Service Act which runs alongside the other EU acts. Next slide.

Here in the United States, we've taken more of a risk-based sectorally-specific approach that's been highly distributed across federal agencies. So some of you may remember that there was an executive order in the last term that required agencies to outline their AI plans, how they were going to procure it, what they were going to do to

regulate, etc.

And now we're seeing that return and there's actually been really great guidance, if I may add, from the Equal Employment Opportunity Commission, who's come up with guidance on hiring algorithms.

I just put this as an exercise that will make our work on this committee quite fun because there are a lot of technical cadence strategies coming out of the various government agencies from CFPB to the FDA to the Consumer Product Safety Commission.

And something I've written about, which is the energy star rating, multi-stakeholder process. There's also a lot of discussion on standards, licensing, consumer disclosures, and now we're seeing an integration of this work with federal privacy. Next slide.

And I promise I am almost done. I think at the end, what I'd like to point out, as I look at this stuff, I was recently at a meeting with the Partnership on AI and a young lady showed this chart that had all these bullets and the United States leaned more towards this voluntary roadmap

compared to our EU counterparts.

When we think about this, we do have the White House Bill of Rights that has been discussed before this committee and we also have the NIST AI risk management framework, which I think are two models that we're currently seeing dominate the United States discussion.

We obviously have enforceable -- enforcement through the FTC, Department of Justice, but for our purposes, I think that these provide some sense and semblance of where the United States wants to head when it comes to comprehensive conversations on regulation and self-regulation when it comes to AI. Next slide.

And I would just say, as every Brookings policymaker would say, what's needed. We always end here. Obviously, less segmentation and more clarity over jurisdictional authority, which is why I think the work that we're going to do in this committee is important, to help us understand this sector even more. Potential harmonization with EU regulation, exploring what's worthwhile to look at.

Obviously, the NTIA request for comment came from I think a domestic with some international implication of harmonization, but I think that's worth further conversation. Are there sectoral regulations by regulated and unregulated industries? When I first started in this work, the first place that I went were to the regulators.

And they gave me guidance on how we should look at emerging models, which goes to this use of generative AI and these call for voluntary commitments. And my last slide, just to close up everything, I thought about, as I was preparing my presentation, just things we should be exploring.

And I'd like to offer to our committee, my colleagues here as well as the Commissioners, obviously understanding the structure and impact of clear sector-specific standards in the area of promoting AI safety.

I think the issue of responsible AI now, based on everything that I showed you, is around responsibility as it plays out in trust and safety, trust and safety on the part of consumers, trust

and safety on the part of the technical cadence of systems. But we should explore that even more.

And the role of either regulation or suggestive guidance in the trading marketplace versus voluntary commitments and really assessing the good, the bad and the ugly of all three on the continuum.

I'd like to ensure that -- to offer up a discussion on the impact of civil rights. To get to more responsible AI, we have to be more inclusive. And I think having the conversation on where we -- actually, the purview of civil rights is important to any conversation that we have around AI.

And obviously the positioning of consumer and industry disclosures, allowing the public interest to be led by transparency. I would say to my friend Travis, I don't think that the box got blacker. It just got more opaque.

And I think it's important for us to continue to append what's behind that so we understand, so consumers can make those decisions on whether or not they want to also voluntarily be part of that

ecosystem.

Unfortunately, I'll end here, I just finished a book on the digital divide that comes out next year. It's not easy because we have not spoken about the people who actually are not online to be a part of any part of this economy. And they are also our customers, whether directly or indirectly, by the products that this Commission serves.

So I'll just put my contact information up. Thank you all for not kicking me off the stage. I feel like I was on the Apollo Theater and I was about to get the ding-ding that your time is up. But look forward to the further discussion on this.

MS. HOUSE: Well, we're all lucky to have a sociologist friend and colleague like you, Nicol, on the TAC.

So at this time, I would like to open up the floor to members of the TAC for questions, comments for our two presenters, Travis, Nicol, and related to their fantastic presentations. If you do have a comment or a question, if you'll just turn your flag over on its side. Corey, I see you're the

first to raise your flag. Over to you.

MR. THEN: Great. Can you hear me? Good. Thank you, Dr. Lee, for really insightful presentation. I enjoyed that. Two questions. The first is, is there like a paradigmatic example that already exists of biases coming out of these models, particularly in the financial context? Since we're at a financial agency.

And then the second is, I would love to hear your thoughts on whether this sort of like dispersed approach that is kind of developing is the right one or whether there ought to be some centralization through a new agency or the like.

MS. TURNER LEE: Yes. Well, thank you for those questions to my colleague here. So it's interesting. When I was asked to join this community by Anthony, I said, okay. This is new. It's a new space.

And what I'm finding more and more, we have seen some examples, Corey, when it comes to financial services based on Latanya Sweeney's work early on, which told us on the financial services

side that people who had names that sounded Black or Latino were offered higher interest, predatory credit card offerings.

My colleague Aaron Klein, and I can share these -- this work with this committee -- said that you're more likely to be denied credit if you're on a PC versus a Mac based on it as a determination of your credit worthiness.

Many of us are familiar with the Apple card where there was a differential impact of pre-authorization, which had a lot of other implications based on who was applying. But it was one of those cases where we actually got to see the product play out with people side by side that were doing this at the same time.

And the GAO report is now suggesting, this just came out in March of 2023 on the financial services side, that we're seeing more discrimination when it comes to things like buy now, pay later and other products that have this assumption that your digital footprint is scaled.

On this side, it's interesting; right? Because

I like to consider, the same way I'd like to look at broadband infrastructure. It's like the backhaul, right, of what keeps markets vibrant and robust. I think the investigation into some prob- -- probable scenarios, without getting too hypothetical, will be an interesting exercise of this committee; right?

And the extent to which we're seeing some gray lines when it comes to algorithmic amplification in trading or whether or not there are other behaviors or protocols that can benefit from AI but at the same time encroach on some malfeasance. So we'll have to go back and look at that more directly.

And for me, that's how the technological space has sort of been evolving where there are consumer-facing products and then there are backhaul products. And AI affects all of them because they're helping us on this efficiency side as well as others.

In terms of the dispersed approach I like to consider this space like a whole bunch of post-it notes on the wall that we're trying to figure out,

who's the right person? What's the right agency? And what should be we be regulating?

And I think Senator Schumer has committed to do a series of talks. We saw the same thing with the White House Bill of Rights. At the end of the day my colleagues Tom Wheeler and Mark McCarthy do suggest a centralized agency for this where you bring in technologists that understand public policy.

We have to go and tell that to universities like Stanford and others who are now bringing up that population of people who live in both spaces.

But it will be interesting to see where we land because I don't know yet if that's the right idea given the fact that, based on the algorithmic violation, we still do have measures in place through the Department of Justice, the FTC, on deceptive practices. It all depends on what part we care about.

And I think again for this committee, that will be an interesting exercise to figure out, what do we care about here? And then, do we have the

structure already in place that we don't have to create a new structure?

Or is it suggest that this actually feeds into a more centralized body, which I think the jury is still out on where that sits.

MS. HOUSE: Vice Chair Redbord.

MR. REDBORD: Thank you so much, really, Travis and Nicol. Excellent presentations. I learned a lot. Nicol, we spoke earlier and you I thought had a really excellent point around the way this technology could potentially be used to empower women and people of color, entrepreneurs.

Could you sort of get into that a little bit? Because I thought that was sort of just a really interesting sort of corollary to the conversation around bias in particular.

MS. TURNER LEE: Yeah. I mean, I had mentioned in the meeting, just to paraphrase, that when we look at societal concerns, we should also think about opportunities.

And so when you think about this particular marketplace in the areas of -- that we're

interested here, what does it look like to use AI to sort of supplement a minority-owned business that may have one or two employees that can actually leverage AI in terms of their capacity to compete in marketplaces that are highly structured, highly bureaucratic?

It's a lower barrier to entry, I believe, in some of those cases, which is something that we're actually seeing with the metaverse, for example, where people have tools, that they can use those tools to enter markets that have been previously contained or where they have been excluded.

And so I think having those conversations, particularly as we look at the growth in the marketplace of Black and women-owned businesses in the space, Hispanic-owned businesses, presenting AI as an opportunity for them to expand and make more efficient their business operations.

We live in a world where, unlike when my grandmother was around, people don't keep their money in the mattress. And we're finding many more people engaging in very diverse practices to get to

the same goal.

And so to our co-chair, that's what I'm thinking. It's often a question that does not get answered because of the perilous nature of this technology, but I do think it's worth a conversation in terms of creating a lower barrier to entry. Yes.

MS. HOUSE: Great. A lot of questions coming up so we'll go Hilary, Todd, Joe and then Steve. Hilary.

MS. ALLEN: Another question for Nicol. Sorry. You're in the hot seat. No. That was a great presentation and I share your concerns about bias, particularly in sort of the provision of consumer financial services

MS. TURNER LEE: Yes.

MS. ALLEN: But one thing I would say about that is at least it's a -- it's a big data situation.

MS. TURNER LEE: Right.

MS. ALLEN: Right? So the work that I've done on AI, I've done it in the context of financial

risk management tools. And the concern I have there is that if we're worried about sort of financial stability, it may not be a big data situation in the sense that we only have one single market history, one line.

MS. TURNER LEE: Right.

MS. ALLEN: And is that enough data to populate these models? And so the -- I guess the big picture question I have for you is, how can we tell if we have enough data for -- in a particular context for the models to be reliable?

MS. TURNER LEE: Yeah. No. I love that question and it kind of goes back to Corey's question in terms of the work of this group.

So I do agree that in the cases of consumer interaction with these technologies, we do have a variety of touch points with consumers, or because we don't often ask them these things, we can actually put together a compilation of various datasets and behaviors online.

What's interesting here is I often tell people one way that we check whether or not we're creating

the same type of scenarios of the data is to, one, check for bias. And that's where I think independent auditing and accountability measures are really important.

But most importantly, in this marketplace that we're working in, where we're relying upon not a variety of interactions but the single source data, is to think about if the data we're using is suggestive of the same structural barriers to entry; right? Into this marketplace.

And I think that's something -- I was just recently telling a bunch of technologists, for me, if your data says the same thing over and over again, it either means that the market has not changed and has become still less inclusive, or we haven't thought differently. That's that design question about how we approach and use these technologies to penetrate markets differently than we have in the past; right?

And so as an academic, that interrogation process at the beginning of this pro- -- of that figure I showed you is really important to me

because it allows us -- I mean, I'm a sociologist. We can find a pathology in anything. You know, you can ask a sociologist about urban poverty. We'll find it; right?

But one of the most interesting things I have found in this space in interacting with technologists is, but what do we do to not find that? And how do we actually create opportunities and space in ways that we can inform it with the existing data but we can also like find other datasets that help us better understand the market in which we want to make those investments, we want to grow.

And for me, that's a question that we often do not ask ourselves because we're trailing for respectability and responsibility by making sure that we're always aligning with the values in human and civil rights laws, statutes, rather than really asking ourselves with this new technology, are we breaking through what we have traditionally defined as our markets?

Which, again, as the only sociologist on here,

happy to answer those questions in our committee meetings. But I think it's just a different way of looking at that question that you just asked.

MS. HOUSE: Great question. Todd, over to you.

MR. CONKLIN: Mm-hmm. Thanks, Carole. I appreciate it. Thank you, everybody. Great to be here. So last session I did brief on Treasury's cloud works. We are going to expand --

MS. TURNER LEE: Yes.

MR. CONKLIN: -- that work to include AI issues going forward. We're still in the planning stages and Steve Silberstein and Kevin Greenfield are all a significant part of -- part of that journey and work to come. But a lot of -- a lot of the discourse in the policy community and in the Beltway area has been focused on the idea of watermarking -- and this is for either of you.

And I think it's in some even draft legislation, a lot of policy discussions. I'm a bit skeptical of that being kind of this be-all end-all, but am I misreading it? How do you all think about watermarking as a potential solution?

MR. HALL: Sure. I'd be happy to. I think that watermarking isn't an end-all be-all by any means. You know, it certainly can help with some of the potential issues; right?

If we're talking about mis-disinformation or attempts to track down mass flooding of information from chatbots, generative AI, things like that, being able to trace it back to a particular model, trying to then correct for the model or being able to identify a mis disinformation campaign by the watermarking of it coming from it.

And of course smaller mis and disinformation campaigns like plagiarism in schools; right? Like things like that, watermarking can potentially help. Similarly, you have the other side of it, which is data provenance, where you can say that a model is being -- is producing or reproducing data and you can trace it back to the actual source; right?

And say, okay. This is actually -- give it -- you know, you ask a chatbot for information about the U.S. population and it's saying, oh. The

provenance of this data is actually coming from U.S. Census as opposed to a hallucination; right? Or another data like some random website or something.

So those can solve for particular risks of particularly these -- the generative AI concerns. It doesn't really address some of the more traditional concerns around straight strict algorithmic decision making in terms of what decisions are being made and why.

It also doesn't necessarily address some of the more systemic risks that Nicol has kind of like talked about in terms of like how it's being deployed, in which manner, things like that.

But for some of the more pressing concerns about generative AI, if we were able to get it to work, right, and there are conversations with NIST and others on technical standards to try to figure out data provenance and watermarking, it could potentially help to mitigate those risks.

MS. TURNER LEE: Yeah. I would agree. I mean, I think this conversation of watermarking has come

up more so, Kevin -- Todd, right, with generative AI because they don't source out where a lot of the evidence is coming as well. And we're not able to track back because of the proprietary concerns that NTIA has found out about.

To me, those are just tools in the toolkit; right? Because I think the ultimate question that we have to answer in this space is, what do we care about as a country; right? Do we care about a bill of rights over this? If that's the case, we will change the behavior in which we manage these technologies.

It's been interesting. The watermarking conversation has come up in my conversations with NIST where they're looking at the risk management framework. You know, obviously they don't have any enforceable capabilities, but it's changing the behavior of how we manage our algorithms and how we manage these systems here.

And I think that's something we can watermark, we can license, we can try to standardize the Substacks. But at the end of the day, we have to

come up with some real key values and guidance on how we embed these technologies in our society.

So I echo everything that Travis said. Just one of those areas that's another tool that -- another Post-it note for me that was thrown on the wall that we're constantly evaluating.

MS. HOUSE: Wonderful. Yes. Yeah. Next, we have Joe and then Steve. Joe, joining us from cyberspace.

MR. SALUZZI: Hi. Thank you. Yes. Thank you for the chance here to talk. I'm sorry I couldn't be there in person. Yesterday, Chair Gensler, SEC Chair Gensler, gave a speech titled "Isaac Newton to AI," which -- really good speech.

But in it he was talking about the challenges of AI and he asked the question whose data is it? And something that we've worked on over here a lot is when it comes to the stock market, because that's mainly my concern, which is a little bit different.

But we talk about the proprietary data feeds that the stock exchanges sell. And that really is a

fuel that drives a lot of the algorithms, particularly these trading algorithms that are out there, which are using AI technology.

But something that came up recently which was more concerning to us was the DTCC, which is the Depository Trading Corporation [sic], they were also selling data feeds. And this was even more concerning because it's a monopoly organization. We only have one place to clear our trades.

They were selling data feeds, which included institutional investor information, which was then segmented to give out even more information. So in other words, this information wasn't out there for the public but if you bought it, it was there.

This caused a big concern in our industry. We actually wrote a paper on it and the institutional community, after numerous phone calls with the DTCC, got them to suspend this product as of the end of this month.

So I guess my question/concern is the data is the fuel that drives a lot of these AI engines. What types of controls can we have over a situation

like this where there -- it is a monopoly, it is proprietary data, that an organization is choosing to use to sell for their own profit, but in the end could be hurting a lot of other investors. Thank you.

MR. HALL: Sure. I am certainly not an expert in that but I can offer some thoughts. I think that the underlying question is one that we are -- that is coming up regularly. It is a question that is coming up in the tur- -- in context of copyright and copyrightability, right, where you have famously a lot of these generative systems are creating images in the style of famous artists.

And these artists are, I think, in terms of a moral sense, perhaps not the legal sense, is still up in the air seriously concerned about the ways in which these systems are taking the underlying information that is copyrighted, has protections, has controls that are in place for how it can and should be used and then essentially breaking through them by, sift -- you know, throwing them into the AI filter.

And I do think that there is a fundamental concern also in terms of potential privacy breaches. And the example that you raised of this information, you can control its sale and access potentially; right? Of the underlying data itself.

But if a model is trained on said data and then the inferences from that data are essentially giving this the -- what you are looking for with that data; right? Then that is something that could be very concerning.

So if you have a model that is trained on proprietary information, insider information, and then that informa- -- that -- not the data itself or the information itself is being bought, but the AI system that can be queried and you can gain inferences from it from that proprietary or insider information. Then that would be something that could potentially be concerning.

Again, I'm not sure if this is -- that there -- that this is something that there isn't some degree of. And again, this is where my area of expertise is -- is limited, where you probably

haven't dealt with things like this before; right? Of like people trying to do a two or three step in order to get around the regulation.

But just simply with the types of systems that we're talking about right now, particularly where that information is being mixed up with a whole bunch of other information so that it is potentially obscured, this could be -- I could see that being an area of concern.

MS. TURNER LEE: Yes. And I would like to just briefly comment. I think we have seen a little bit of this, not in this particular area, but in facial recognition technologies, where we see FBI, mugshot databases, driver's license databases, being mined by third parties that use what they collect to actually place them into proprietary databases to be used for facial recognition scanning at airports, at prisons wherever the case may be.

But to your point, being used outside of the organization or in- -- individual's consent and being sold in the third-party market, in addition to being leveraged by government in ways that it

has some potential civil rights violations.

And that's something to me that is really interesting in this conversation, because I do think kind of goes back to Hilary's question; right? So what do you do when you have one single source data and you have a monopoly of data brokers that actually have that data and make determinations on what they want to do with that?

Obviously, federal privacy legislation of some form could help in terms of the collection, retention, storing, repurposing of institutional as well as individual data.

But some guardrails and guidance from regulators who actually are in this space as well that goes back to whether it's a voluntary commitment or if it's hard regulation, to actually make sure that they're adhering to some kind of standard when it comes to data handling is also important as well.

So again, I would -- I would lean back, Travis, honestly, on the FRT databases that are being developed without a lot of people's knowledge

and particularly added to that some of the newer social media surveillance of photos that is actually being used for proprietary purposes and how that may also apply to this case where we have proprietary institutional information being resold on the third-party market that has implications going forward.

MS. HOUSE: Thank you. Over to Steve and then we'll close out the discussion with a question from the Commissioner. So Steve, you're up.

MR. SUPPAN: Yeah. I had a general question about, how do AI models train on data gaps? And I have in the background of my head that the use of AI for surveillance of position limits and position aggregation, if you have either exchange rules or federal rules which exclude what I would regard as relative data elements that should be reported.

MS. TURNER LEE: Mm-hmm.

MR. SUPPAN: How -- how does AI enable surveillance of data gaps if it -- if it does?

MS. TURNER LEE: Mm-hmm. You know, I would say that that's an interesting question and thank you

for bringing that up. I mean, I think it does control for data gaps because, whether or not your information is included or excluded, it actually permits activity in these spaces that we're talking about.

A great example there, ride sharing. If you are a person -- there's been studies from MIT and others. If you're not part of the ride sharing community, just as recent as a few years ago on the other side of the water in Anacostia, people weren't using Uber.

Guess what? They were not part of the ride sharing ecosystem in two ways. Uber didn't go there and Uber didn't go there; right? Because they didn't know about the people and then they didn't see them as a viable marketplace.

So the system does control for those data gaps. From a policy perspective, some of the areas that people like myself have been advocating for in the financial service space as well are safe harbors or sandboxes to help us identify areas where we need to fill in those gaps.

So when we're trying to put out -- we saw this actually in fintech where we wanted to get more information about what could potentially create greater bias. And we did some co-evolution between fintech companies, consumers, organizations, etc. We should see more of that, I think, to help us identify those gaps.

The challenge is tech companies do not collect information, my last point on these gaps, because by law they cannot. So they have to make assumptions, which then goes back to those inferences.

Unlike a researcher who has to take care in human subject handling, that's not always the case in industry. And they make propositions and suppositions, presumptions around those populations that will lend them being part of the ecosystem or excluded from the ecosystem, which I think is again another area that we can discuss as well here.

MR. HALL: So not being somebody who's deeply technical, I do think that there's -- there is a certain degree of promise and/or peril, depending

on how you're looking at it, in these technologies in that, like other types of statistical inference models, they are able to infer what they do not know; right?

They are able to themselves fill in some of those gaps. And sometimes with unnerving accuracy and sometimes with unnerving inaccuracy. And so I think that there is -- I mean, again, like there's a degree to which these mo- -- you know, particularly the frontier models, are promising in that regard, if you want to fill in those gaps.

And but of course there is some degree of risk with using them in that way, or particularly where there are significant data gaps, where you do end up creating systems that do become -- and Nicol referenced this before, right [inaudible] parrots, right, who recreate systemic biases in the ga- -- in the data that is already exists; right? So it just simply makes it even more reified.

MS. HOUSE: Commissioner, close us out.

MS. GOLDSMITH ROMERO: First of all, this was incredible. Thank you, Travis, for coming over here

and sharing your work. Thank you to Dr. Turner Lee. You've given us a lot to think about and we will roll up our sleeves and continue to work on this.

And thank you for all of you who are participating. I mean, this is really -- I think we're -- we're learning more and we're getting deeper into it. My like sort of closeout question relates to governance. And when you look -- first of all, I -- as I said, there are many generations of AI being used for financial services.

Obviously, when anyone has sort of played with ChatGPT 3 or 4, often you see it sort of not laying out both sides and maybe not making a choice or make -- not making judgment. But when you -- when you hear from Dr. Turner Lee, there is clearly judgment and choice going into the questions, going into the way the model is created and the outcomes.

And so my first question, Travis, is, in NTIA's work on governance, what are the types of things that might come out of your request for comment? And then also to Dr. Turner Lee, are there best practices on governance out there, other

things? I know you were sort of talking about some of the hard and soft laws.

MR. HALL: Sure. Thanks. So I mean, our final ultimate product is going to be the report that we're issuing in the fall. I would also say that there -- you know, bleeding into the other question, there's -- we are actively participating in interagency processes that are looking to address a wide scale of risks that are coming with these models and potential actions on that.

And certainly the administration has been actively engaged internally in these processes, right, in terms of 2019, trying to figure out its governance models. And I think that the lesson from that is that it's really hard; right?

I mean, the U.S. government that should have records of the things that it procures and how it's being used and things like that, in terms of -- the definitions are extraordinarily important in terms of getting the definitions concise, applicable and understandable in order for these agencies to understand what is being asked of them in term --

in terms of what -- how they use these models and what their governance structures are.

And so I think that what we will come out with in terms of the -- at the most important part of our report will be some recommendations around policies for specifically auditing and impact assessments and certifications and other things like that, that will likely also apply to the actual governance structures within companies. I will say that for this group, what might be of interest is that investors are extraordinarily interested in our work as well.

They are extraordinarily interested in this because they are similarly confounded somewhat by investing in companies that are developing these tools, knowing that they are actually producing or going or have the ability to produce something that is a meaningful service and also something that is not going to be particularly harmful or -- and ultimately a bad investment.

So we've heard from multiple investors groups ab- -- of interest in accountability for the

companies that they're investing in. Thank you.

MS. TURNER LEE: Yeah. That was a great point, Travis. Thanks for sharing that too. Forgot about that. I would say also, Commissioner, three things. I'll just end here. Obviously, consistent multi-stakeholder input is something that we're seeing on the governance side that's working really well.

It's working quite well overseas when it comes to multi-stakeholder processes. We're seeing these in the UNOECD, but what's happening is we're bringing various actors to the table to speak about that.

And we tried a little bit of that with the White House Bill of Rights through listening sessions and public feedback. So the more we're able to glean public feedback and give opportunities for people to come from all walks of life to speak about this is important. Particularly, having investors at the table also matters. So I'm glad to be a part of this.

The other thing that I think has been a best practices guidance, the chairwoman of the EEOC has

been the first to really lead, I think, with some great guidance on how we look at AI in hiring, how we look at AI in compliance to ADA requirements. And the GAO is about to come out with something similar which I think are really good pieces of literature to follow.

May not necessarily be enforceable, but it allows for great guidance on various government ver- -- sectoral verticals. And then the last thing I would say is education. We often miss this side of this debate which is educating people about this technology, both its opportunities and its perils.

And so I think what has worked that I've seen is much more AI literacy which often is an afterthought at the beginning so that people understand how these systems work, where they stand among the continuum of these systems, but more importantly, where it stands sectorally. So having more of those conversations as well.

MS. HOUSE: Thank you all so much. Echoing the Commissioner's comments on what a great discussion.

So for our third presentation regarding AI

issues, which very helpfully intersects also with our third topic and later presentations regarding cybersecurity, Dan Guido, co-founder and CEO of Trail of Bits will be presenting an impact assessment of the proliferation of AI cybersecurity capabilities on financial security. Dan, over to you.

MR. GUIDO: Thanks, Carole. So as if we weren't all concerned enough about AI, I wanted to throw a curveball. So Trail of Bits is a cybersecurity research and development firm. We help companies identify, exploit and mitigate cybersecurity risks within their firms.

And we specialize in emerging technology, so there are really very few other people in the world that are as capable as us as identifying how hackers can hack better with AI and what systemic risks that might have on the financial security industry, on the finance industry.

So from my perspective, the availability of these models should cause us to reassess the defensive techniques that we apply because the cost

for attackers to perform certain attacks will change dramatically with the availability of these techniques.

Normally attack and defense, they're -- they're somewhat static. We have an equilibrium. We sort of know what exists and it's a -- it's a defined number of defenses we need to implement.

But we're entering a period of much more extreme turbulence where we don't actually know what these models are capable of. And there are some initial results that it may actually dramatically affect the ability of hackers to be effective.

So slide t- -- oh. We're way -- we're way beyond where I -- go -- go back. My slides are off. Thanks. Slide 2 is great. Okay. Amazing.

So from the outsider's perspective or from a known expert's perspective, it can be really hard to figure out what's accurate. You have a lot of proponents that are saying that ush- -- AI is going to usher in some kind of utopia or that it's going to kill us all.

On the other hand, you've got a lot of people that are claiming that this is just big tech puffery, that there is nothing here, it lies all the time, it can't produce any sort of realistic results for anybody, that's it a giant waste of time.

We've seen this in the security industry as well. You've got security researchers that plug a piece of code into it like the screenshot, they drop a theorem contract into -- in it, and it immediately found a security vulnerability.

On the other hand, you've got a consulting firm that says, "Oh, we tried that for ChatGPT, doesn't work at all, completely misdirects you, you shouldn't use it, period."

And sometimes you've got companies that are both saying one thing and the other at the same time, like OpenAI. The thing is that a lot of these people, first off, aren't evaluating the technology in the way that I think we should.

On -- on the left here, one of those proponents didn't realize that the data they were

testing it against, the theorem contract with the vulnerability unit, was in the training dataset, that all they were asking was the AI to regurgitate information that it had previously memorized.

On the other hand, I think a lot of people are looking at AI to be magic like some Jedi mind trick that can immediately discover a security vulnerability with the simplest of prompts, but that's also not how this will work. We -- we already have general human intelligen- -- we -- we already have general artificial intel- -- or -- artificial -- we have general intelligence. We have humans.

Humans can be accelerated dramatically by taking good advantage of these sorts of techniques. What is a mid-level -- is it possible for a mid-level engineer to operate as an expert with the addition of AI techniques? That's really the question that I want to understand.

So next slide. At Trail of Bits we've tried to do this. So I have some empirical results from experiments that we have run internally to see how

this will help our ability to identify, exploit and mitigate security vulnerabilities.

We need to realize a few things first. We have to understand that AI is not magic. It's just another tool in our toolset and that we have to apply it to the right problems. We had to learn how to use the technology. We couldn't just ask it, "Please find the bug for me." It's not magic. You have to figure out things like chain of thought prompting and ex- -- asking an AI to explain your reasoning, or a lot of the other prompt engineering techniques that I think people have -- have now started to discuss.

We also have to evaluate progress correctly. We need a benchmark against which we can decide if this is actually helping us or not, or again, many different tools that we can choose to use, is this better than the state of the art, is there something else I can use that's algorithmic in nature that instantly gets me the answer, what's the bar of the progress.

And then we have to choose the right problems,

because AI think a lot of people probably try to ask ChatGPT for help with their math homework, and they realize that it's really not good at that, because that requires correctness.

But there are lots of problems that security researchers work on that don't require correctness, that require breadth of knowledge, that require -- that -- where -- where mistakes are acceptable.

So with that, we've been able to use to decompile code into high-level languages that we could never do before. We've get it -- we've gotten it to identifying trigger, software security vulnerabilities that are present in browsers.

We can use it to greatly accelerate the offense of capabilities of a mid-level cybersecurity engineer. So let's go to next slide.

So from a sort of anecdotal perspective, this is how we think about it, where we've evaluated the efficacy of AI in the life cycle of a cybersecurity engineer. We look at problems that require breadth and where mistakes are acceptable as the best opportunities to first apply it; things like

documenting a function. Well, there's probably no documentation for it before. You can use that to audit a code base more quickly.

If you have a code base of 10 million lines of code and you want to very quickly understand some broad information about it, AI is phenomenal, because even if it makes a mistake it at least points you in the right direction.

Phishing emails, another really great one, which I'll tell a funny store about in a minute. But I think a lot of people realize that the ability to develop, to write human language at speed and scale, is something that is a brand new capability for the world that didn't exist before these were available, before AI models were widely available.

So there are no longer sorts of typos that you can watch out for. You're not going to get a generic message; you'll get an extremely targeted one. But things that it can't do are things that require mastery or where perfection is required.

I think an analogy would be that you wouldn't

really ask AI to paint a Sistine Chapel for you, but it's going to be able to create an endless stream of clipart for every presentation you'll ever make, and it can do that 24/7 without taking breaks.

So yeah, I think one -- another analogy is it's kind of, imagine if you had an infinitely scalable team of extremely sleep-deprived graduate students, it's -- you know, you're not going to get perfection, you're not going to get mastery, but you will be able to get some results. So next slide.

So based on that somewhat anecdotal evaluation, from a cybersecurity perspective, there are a number of technologies that are now at extreme risk of being obsolete because of this -- because of these accelerants that have been applied to attacker behavior.

Chief among them we have great evidence already for bug bounties and phishing training, encountering significant difficulty with the -- with today's availability of AI models.

The incentives -- bug bounties align perfectly that -- for Third World developing countries, the bounties that are available $500, $1,000, are an extremely meaningful amount of money to many people.

So it makes sense from their perspective to spam bounty submissions to whoever is willing to receive them. And because they're produced with an LLM, they sound reasonable, they sound confident, they sound like they could be accurate.

And they might be to like an 80 percent degree, but only a highly paid expert can unravel the mystery of whether or not they represent a real vulnerability. And that mismatch in incentives is going to cause significant problems with the efficacy of these programs.

Same thing for phishing training. ChatGPT knows who I am, right, it knows who many of the people in the room are. It's trained on the entire internet's dataset and you can ask it to write a perfect phishing message to reach me, and it turns out an extremely good result.

You can ask it for 10 different copies and send me a new one every week until I get -- until I get fooled. So this really implies that for all these things, for bug bounties, phishing trainings, signature-based defenses like iDS or anti-virus as well as threat actor attribution. All of these were attempts to deal with the problem without solving the underlying issue. Right?

We have code that's insecure. We have computers that when you hack them, you can take over a whole network. You know, we have to identify -- we have to actually address some of these core cybersecurity problems without fixing its effects, with defenses like these. Next slide.

So as a regulatory agency, one thing that you might ask yourself is how can we regulate this problem away, are there -- are there methods that we can apply to these AI models to force them not to produce these sort of outcomes? And the answer is absolutely not.

This is not an opportunity. Alignment will not save you. You cannot reinforcement, learn your way

out of the problem, and you can't restrict availability of AI models.

As of 12 hours ago, the best open-source model that was available was completely conceived of, researched, developed, and then made available by a firm based in the UAE, completely outside the United States. That's Falcon 40B. Right?

And these models that are available, that are open-source, that are completely open to being fine tuned by anybody for anything, are quite good. They -- they start to compete the a- -- availability of models like GPT4 or Claude, etc.

Now, these models, because they are so effective, because they've been trained on so much data, always within that dataset, are going to be the knowledge how to commit crimes, of how to hack into things. And you can't unlearn that from the dataset.

Output censorship is really an unwinnable task that just frustrates amateurs and researchers. There's a really great sort of example or reusable jailbreak, people call it, called DAN, Do Anything

Now, which is a set of text you can drop in to a bunch of different LLMs to get it to just invert its understanding of what it should do and should not.

And this is also a really surprising result. But the more that you try to align an LLM, the more that you try to tell it what not to do, you're actually helping it build a model of that bad behavior, so that at some point in the future you can invert that understanding. And they're very quickly very good at the thing that you try to get not to do. People in the community call that the "Waluigi Effect."

Let's go next line. Now, all this is great and it's based on a lot empirical data that Trail of Bits has invested in and has performed on our own. But what we don't have is we don't have any kind of systematic measurements of what these models are capable of doing right now.

A lot of the effort in this space is towards software programming, it's towards software development, how good is this thing at writing

JavaScript or Python. If I have to change a form on my website, can it do it or will it break it.

And these evaluation benchmarks are quite rigorous. They do -- they have thousands and thousands of tests. They are aware of specifically the information in those tests. There's just lots of capabilities and features that -- things like HumanEval or EvalPlus enable for us to actually measure progress of models like GPT4 or Claude at writing code.

However, we have no such taxonomy and no such benchmark for cybersecurity. We don't know if these models can solve one problem out of a thousand that I deal with every day, and we don't know how that compares to the state-of-the-art tools that I have available already. And we don't know how this might affect the daily work of a junior, midlevel, or expert level cybersecurity engineer or hacker.

So this is a major research area for Trail of Bits, but it also needs to be a community resource. It needs to be something that the financial community understands so they can identify

objectively whether their defenses are at greater risk are now made obsolete by the availability of these models.

So yeah, we know there's stuff it can't do; it's not magic, but there is a bunch of stuff it can really help with. And we just don't even know because there's no systematic evaluation. So next slide.

So again, going back to hype and the advertising I think that there are right ways and wrong ways to build these systematic evaluation models. Right? I think a lot of the ways that people evaluate LLMs right now is they ask it to pass, for instance, exams like the LSAT, the MCAT, the MIT final exams.

But first off, a lot of these evaluations have trouble because that data is already in the training data site. You're just asking the LLM to regurgitate stuff that it memorized, which anybody can do.

So you really want to know, does the output -- like what happens when you give it new problems,

novel problems. And in order to create an evaluation metric that operates that way, requires a lot of investment.

You can't do these really cheap and simple things like download a bunch of final exams from MIT's website and pump them through LLM.

You also want to know if it can generalize past a few examples. You can very easily prompt hack your way into getting an LLM to answer one thing for you. But is that going to work time after time after time that I do it?

And then finally, you have to deal with the fact that these things makes mistakes. They hallucinate, they are not 100 percent accurate, they are probabilistic, they are statistical, they're not -- they don't have actual knowledge of what you're asking it to do.

So are mistakes acceptable and will the task that you're asking it to do break, because it does make a mistake down the line. So again, repeating the call, we need more measurement, we need task-specific datasets in the evaluation frameworks,

especially for cybersecurity.

The sorts of questions that I ask it like if I want to point it in the right direction or is this technology, even if it's sort of feeble, remove a serious constraint that my adversaries or I had.

So for example, crafting phishing languages in any language or at any person of any sort of social class or role within a company, how rapidly is it improving, is something that I would like to know.

So in the past, techniques around code synthesis were terrible and now they are much less so. But at what point does it exists right now and how fast might we get there. And is there sort of a synergy with existing technology or tools.

For example, what does the extreme pressure and the availability of GPUs have on adjacent technologies or other sorts of problems in cybersecurity. Right now we're producing GPUs at an extremely fast rate. That field is advancing more so than it has in the last 10 years.

And lots of companies have these giant clusters of tens of thousands of them; are there

other tasks that can be solved or that might be affected by the availability -- the easy availability of extreme levels of compute.

And then finally, our mistakes okay would be another -- or easy to catch, would be another metric that I would use to figure out if there's some fundamental change. Okay, next slide.

So to wrap up here, AI is a systemic direct of cybersecurity. It changes the cost model for attackers and defenders. A lot of the ways that people are looking at AI to change cybersecurity is by magic. They think that it will be a Jedi mind trick, but it's not. It's an augment human -- human capability. It's going to remove constraints that attackers previously had that they no longer do.

And if you try to fix this problem by leaning into alignment or control of model availability, then you're not going to be spending your time effectively. We need to create measurement and benchmarks that enable us to figure out how at risk we are and also what opportunities that we have.

So from a positive note, when I look at attack

and defense on cybersecurity, defense has a lot more problems that look like I have an infinite amount of data to sort through; I need to -- I need some low-trained person to review it all, whereas offense has a lot more problems that look like I need to paint something like Michelangelo.

Offense you need to construct these articulate attacks that work, that are interpreted by a computer on the other end correctly in order to achieve the outcome that you want.

So while I think the short-term is definitely more geared towards offense, the medium to long-term may be more geared towards defense. But in order to get there, in order to build defensive technologies that use AI, we have to be able to experiment with them.

So a trend that I see among a lot of banks is that they've just made blanket restrictions on the ability to use a lot of these models and wanted to shut their head in the sand. But it's actually more important for us now to figure out how we can use these to aid our defense in order to keep up with

attacks from the other side.

So Trail of Bits has a number of resources that are available that directly address this topic. We have our published results from our trials of using AI to beat humans in auditing code for security vulnerabilities, a curative list of references that we believe help someone get up to speed on the field quickly.

We also have our comments we made to the White House OTSP on how AI might affect national security and those sorts of problems.

And then finally, we do also consult with people on how to improve the security of AI-based systems and increase the safety and security of new technologies that they have built as well as techniques for doing so, in that how to measure safety of AI-based systems link.

But thank you. That's my talk.

MS. GOLDSMITH ROMERO: Thank you so much, Dan. I appreciate it coming from our cyber background as well as a few other people sitting around this table. Really appreciate you examining the

intersection of these two key priorities on one of the three major areas for the attack and given the significant issues.

So I'm just opening up the discussion with my own comment that I know that in the world of cybersecurity, there's just generally been issues with sufficiently measuring and understanding cyber risks, which led to issues like mis-calibration of the cyber insurance sector, certainly continued challenges and debates about what the right metric should be for key leaders to decide whether a set resources and defenses and then also being a big sci-fi fan and coming from the military or thinking about what the future of warfare looks like and that actually there were constant -- like cyber warfare will just be perpetual forever and will be AIs set against each other.

This is the stuff that keeps me up at night. So thank you so much for the presentation and scaring us all but also highlighting that there's a way forward, if only we can get some of these steps right, so to focus on the defensive capability is

to be able to mitigate the risks presented by the offensive augmentation by AI.

So thank you, Dan. I'd love to open to the rest of the TAC for their comments. I see Justin. You've got one, you can kick us off with a question.

MR. SLAUGHTER: Thanks for that, Carole. Great presentation. I guess I'm curious about why you're confident that this benefits the defense in the medium term. My understanding, and I'm not the expert that some of you are, is that cybersecurity has favored the attacker basically since we had wide use of the internet, that no one's really found a way for the defense mindset to catch up.

And this strikes me as something where, based on your own example, attackers can make a lot of mistakes because you're fine with getting a very low output of success. If small number of phishing attempts, small number hacks have high value, well, defense needs to be nearly perfect.

So what is the idea here that potentially could shift AI to using where defense is

benefiting?

MR. GUIDO: Yeah. Sure. So at whole attackers only have to -- defense has to win all the time and attackers only need to win once. I sort of think about that a little bit as a -- of a myth.

Attackers actually have to enter a very treacherous environment once they break into a company. They're in a network that they've never observed before, they don't know what monitoring they're subject to, they don't know what's real and what's not. There could be canaries or honey pots, or things that are put there to strategically to distract or catch them.

They're -- and then the tools that they work on are usually very expensive, articulate, require a lot of time. They're the Michelangelo. Right? And those tools are very brittle, because once the community is aware of what those tools are, if you report them to the offender that exploits a zero day, those tools go away, they disappear. They're instantly devalued.

So from an attacker perspective, there are

lots of different levers that you can put effort on in order to make their life hell. And with defense, I think the nature of defense is one where you have to think about breadth all the time.

And it's really difficult for a human to pay attention to everything that's moving around the entire company, all the different datasets, all the different continuously streaming logs, all the different complex architectures of applications.

And AI can find a way to synthesize that and provide me an up-to-date picture all the time, in a way that's nuance, that has context. So I think that just because the nature of those two challenges, there may be ways to use AI much more effectively on defense to change this back into their hands after an initial period of what's going to be pin.

MR. SLAUGHTER: It sounds like what you're saying is that AI allows for redundancy be dynamic, in a way that's not the case right now.

MR. GUIDO: Can you repeat that?

MR. SLAUGHTER: It sounds like what you're

saying is the AI can allow systems to be redundant -- defenses could be -- to be redundant and dynamic in a positive way, that's not true right now?

MR. GUIDO: Yeah. Definitely much more dynamism, definitely, you know -- it's really like, again, what could I do if I had an infinite number of sleep-deprived grad students.

By the infinite number of mediocre junior cybersecurity engineers, what would I be able to look at, what would I be able to inspect, what would I be able to review in real time that today I can't, today that I have to have a batch process that only runs overnight or that I have to have a quarterly meeting about.

So I really do think that there's some opportunities here for defense. But again, in order to do this, we need to have a systematic evaluation of what these sorts of things can do and anthropological understanding of what happens in the day of a cybersecurity engineer, what are all the technology and techniques that they employ, and how does AI relate to each of them, and could it

overcome constraints that they have.

And right now we don't have that for cybersecurity but we do for lots of programming. So we need to build analogous frameworks if we want to ever get there and achieve the upside.

MS. GOLDSMITH ROMERO: Incisive comment. Nicol, next.

MS. TURNER LEE: Here you go. Thank you so much, Dan for that presentation, and appreciated just the conversation of cybersecurity and AI in that intersection.

I do have a question on how do we keep these systems updated? Because what we also know is that AI, once people are sort of onto the AI, change the AI, so that we're always evading those types of cybersecurity mitigations, strategies you discussed.

So just curious, is it documenting? Is it being a step ahead? How do we do that when we as policy makers and innovators are always behind?

MR. GUIDO: Sure. So I'm going to answer from a technology perspective. I mean, we need to fix

the root cause problems. Like phishing education was just a Band-Aid on top of the root cause problem that when somebody in your company gets hacked, it can bring down the entire network that you own.

That should not be the case, it should never have been the case. You know, authentication should have been strongly authenticated, there should be limited information available on that person's machine.

You should have authentication by -- or authorization by context, you know. You need to have an incident response procedure or a minimized sort of blast radius of when a hack happens, what is the worst thing that can occur inside the company.

It shouldn't be that I can hack an executive assistant at Coke and then steal the formula for Coke. Right? Like, these things should be isolated and your company should be safe no matter what emails people click on or what documents that they open.

So that's the thing that I think we need to get back to. I think that we need to really ask ourselves these hard questions again of, are we just sort of distracting ourselves from the root cause problem, are we putting a Band Aid on the issue, or have we ultimately addressed this risk in some way.

MS. HOUSE: Thank you. If there's no other questions, what a way to wrap up a fantastic discussion of two of the three key priority areas. Again, thank you so much, Dan, for the great presentation and for the really insightful questions. So now, we'll take a 10-minute break, and we'll reconvene at 2:15.

MR. Biagioli: Reminder to turn off your microphones.

[break]

MR. REDBORD: Welcome back everybody.

In our last meeting, we heard presentations and engaged in discussion on the challenges and opportunities in a more decentralized financial system. We discussed illicit finance and national

security risks and how to mitigate them through the use of blockchain intelligence and innovative technologies, such as digital identity.

Today we will continue that discussion with a deep dive into regulation and governance. To begin the discussion, our first presenter, Anthony Biagioli, special counsel to the director of the Division of Enforcement to CFTC, will kick us off with an enforcement case study on Ooki DAO. Tony?

MR. BIAGIOLI: Thanks very much, Ari. I want to thank Commissioner Goldsmith Romero for the opportunity to present today. I should note that my remarks today reflect my own views, are not necessarily the views of the CFTC, any CFTC Commissioner or the Division of Enforcement.

The CFTC's litigation against Ooki DAO, a decentralized autonomous organization, arose from activities presenting a fundamental and explicit challenge to the ability of U.S. regulators to enforce the law against groups of decentralized actors using novel, smart contract based technology on blockchains that enable collective decision-

making.

In short, the question is, can a DAO act and be sued and served and ultimately held liable as an entity in its own name responsible for the actions carried out in its name.

Or are DAOs immune from suit, from service and from ultimate liability leaving only pseudonymous individual members susceptible to liability for a DAO's conduct. This isn't just a question from my perspective about government enforcement. Imagine an individual -- a retail individual defrauded by a DAO. What recourse does that person have?

Can that person sue a DAO, recover against the DAO's treasury assets or is that defrauded individual left to pursue pseudonymous individuals that may be difficult to identify and may be located anywhere in the world?

What if the DAO itself is wronged? Can it sue in its own name or is it left once again for individuals with questionable standing to attempt to pursue lawsuits on the DAO's behalf?

The Ooki DAO originated as an LLC doing

business as bZeroX. That LLC operated the trading platform, which using a smart contract-based protocol on the Ethereum blockchain, enabled members of the public to make heavily leveraged bets on the comparative price performance between two digital assets.

Other members of the public could supply liquidity to the protocol, digital assets the traders could borrow to establish their leveraged position.

And I've put on the slide, a heavily simplified sample reflecting the mechanics of a 5x long ETH versus DAI trader opening a leveraged position on what was originally the bZx protocol and later labeled the Ooki protocol.

And essentially, if I'm a trader and I'm real bullish on ETH, I can and I want to leverage up that bet, I can send collateral to the token smart contract, which in this case would borrow the stablecoin DAI from liquidity pools that were funded by members of the public.

The smart contract would send that stablecoin

to an on-chain decentralized exchange where the stablecoin would be swapped into ETH.

The swapped ETH would be sent back to the smart contract and a tokenized position reflecting the trader's 5x long position would be created and provided to the individual. That token could be redeemed at any time. If you bet right, it's redeemed for a profit.

If you bet long, if you bet wrong, then the over-collateralized position, the collateral that you supplied, would be liquidated by mechanisms contained in the smart contract.

I'd pause for a moment to note one of the remarkable features of this enforcement action was how unremarkable the transactions themselves were. I think there was general agreement that these were leveraged, retail commodity transactions that could only be offered to the general public on registered exchanges.

The novelty rather was in the organizations created to offer them. In 2021, the LLC transformed itself into a DAO and transferred operational

control of the protocol to a DAO.

Liquidity providers and others who had received governance tokens entitling holders to propose and vote on any question relevant to the operation of the protocol, were now in control.

Notably in creating the DAO, a founder of the original LLC presented what I think was a fundamental challenge to regulators everywhere in fairly stark terms, where the founder told DAO community members on a community call that it's really exciting.

We're going to be preparing for a new regulatory environment by ensuring that bZx is future-proof. So many people across the industry right now are getting legal notices, and lawmakers are trying to decide whether they want DeFi companies to register as virtual asset service providers or not.

And really, what we're going to do is take all the steps possible to make sure that when regulators ask us to comply, that we have nothing we can really do because we've given it all to the

community.

So from my perspective, the ultimate question presented by this litigation was whether that founder was right. Does transfer of control of a trading protocol from an LLC to a decentralized group immunize that group from obligations to comply with the law simply because they operated in a decentralized manner?

These questions were litigated in the Northern District of California and the Court made three significant holdings. First, you can sue a DAO. The CFTC's position was that a DAO is an unincorporated association that two or more people acting without a charter pursuing a common objective, and the Court held that the DAO was an unincorporated association who could be sued.

Second, you can serve a DAO. By their nature, DAOs lack many of the features that are typically considered predicates for service under federal law and many applicable state rules. They lack many of the characteristics that most service rules presuppose that an entity will have, a physical

presence, a C-suite, an agent authorized to accept service of process.

When you have no physical presence, you're an online business only, you are characterized only by the decentralized collective decision-making of your members and no one, presumably, at least to the public, is in charge, then who is there to serve?

Most service rules again, presuppose that someone like that exists in the organization. Think like the old Elks Lodge, right, an unincorporated association.

There's usually a president and you can serve that president or you can serve by mail to their physical location. Not so in the case of many DAOs.

However, under Ninth Circuit precedent upholding service via email, the Court held that the CFTC service on the DAO via a help chat box on its website with contemporaneous notice to the DAO through its online discussion forum, which was the only mechanism that the DAO held out to the public to contact it, satisfied applicable service rules

and constitutional requirements.

Finally, the DAO is a person under the Commodity Exchange Act. Only a person can violate the provisions of the Act that we charged in this litigation. And the definition -- a person is a defined term.

The definition of person includes associations and the Court held that the DAO is an unincorporated association, not just for purposes of capacity and service, but also for purposes of substantive liability under federal laws, like the Commodity Exchange Act. So the Court held that it's an unincorporated association.

That's a type of association; it's thus a person and the kind of entity that you can sue. So if you're an LLC who is operating a trading protocol, and if that trading protocol is operating in an unlawful manner, from my perspective, this case stands for the proposition that you can't simply transform yourself into a DAO, continue to offer the exact same unlawful trading platform without accountability.

So going forward, there are all sorts of novel and interesting ways that DAOs may seek to contribute to decentralized finance ecosystems. And again I think this case stands for the proposition that when they do so, they do need to comply with the law and can be held responsible if they do not. Thank you.

MR. REDBORD: Tony, thank you so much. Before we dive into a DeFi discussion, we're going to do a second presentation, this one from Ben Milne, founder and CEO of Brale, and Justin Slaughter, policy director at Paradigm. Ben and Justin will jointly present on the extent of decentralization and models of governance in DeFi. Gentlemen?

MR. SLAUGHTER: Thank you very much, Ari, and thank you Tony for that. That was very illuminating. So I drew the shorter straw. So I'm going first. So you can I'll try to get through this quickly so we can listen to Ben's greater wisdom.

So decentralization and DeFi governance, thank you for the chance to present on this. I actually

think this is perhaps the most interesting topic in a sea of interesting topics in crypto to me.

Now, I wanted to kind of as a brief reminder for everybody, DeFi is even in a space as new as crypto, especially new. The space really began about five years ago, with the Uniswap protocol version one white paper, which is 2018 release, and then blew up in DeFi summer in 2020.

I would also stress the most notable thing about DeFi in the last 18 months is actually the degree to which it withheld itself during periods of significant crypto market stress and traditional financial market stress in both 2022 and then this March during the banking issues and Silicon Valley Bank and other banks.

Next slide, please. So there -- we view at Paradigm that there are four key concepts that can in many ways, encapsulate DeFi. The first is self-custody. That's simply the act of holding crypto, and you think about this -- that is in many ways requisite for DeFi.

Because if you have a custodial entity

engaging in crypto financial activity, that almost by definition is not decentralized. There's a central activity. Second, it's autonomous. That means it does not require humans to approve individual transactions.

Again, that does not necessarily mean there's no human involvement. But if you have a DeFi protocol that's being directly approved by humans, it probably does not merit the name DeFi.

Third, transparent. This is fully visible on-chain for the same reason you might expect. If you have a lot of off-chain transactions, it's not a true decentralized financial protocol because, again, that's another black box. The idea of DeFi fundamentally is a minimal amount of black boxes fully visible to everyone on-chain.

And then fourth, that it's interoperable and composable. That means you can first exchange data with other applications; and second, you can build applications on top of each other.

These are the Dapps, basically, that make up much of crypto's Web3 ecosystem. Now, the really

important thing to grasp about this is that decentralization is a spectrum, it's not a toggle. There is no one way to do DeFi or decentralization.

You think about it. On one end of the spectrum, you have something like EDX, which is a recent centralized crypto exchange comprised of several companies working together. They have custody; they fully have a black box.

They entirely control the system. The other end would almost be pure software open source code like Linux. Nobody has any control over that. It exists in the universe. It can be used by anybody.

There is no one entity, not even a lab, and then between that you have a whole host of other entities from any number of DAOs, all of which are structured differently, to any number of labs. So that's the most important here to keep in mind. Next.

So you can see here, this is Uniswap's protocol and its cumulative volume since 2020 basically. You can in many ways see -- and forgive me for the chart crime here because it's cumulative

-- the large linear increase in activity over time.

Now, this chart ends in May 2022. My understanding is there's been continued growth since then from talking to members of the community and this gets at the number one problem we have, which is we still have limited data and research into how this space actually works.

That's a function both of a lack of focus on DeFi until recently and it's also to some extent a regulatory failure. There is still, because of a lack of certainty about how the regulatory regime will work for DeFi, not a good way of getting the information in question.

More than anything else, this is something we need to really work on as we consider regulatory solutions in DeFi, is getting more data about how it operates. Next slide.

Now this comes to DeFi governance and there is in many ways, a natural tension between heavy governance and DeFi.

Now that in fact makes sense because our view is the core virtue of DeFi is credible neutrality.

The idea that no one click, no one entity or group consistently has control or ever has control even over how the protocol works.

Now, the reason for that is fundamentally if someone has control of the protocol, can change it at will, it ceases to be decentralized. At that point, it has become even if it is nominally decentralized, a centralized entity because of individual or group control.

That means that DeFi depends upon first dependability that, you know how it's going to operate, that you don't get surprised because of how it operates, and second that act of avoiding capture. Now, to some extent, the greater the governance or rather the more there is, the greater the risk of capture.

And that's because the more protocols, the more requirements you put in place, the easier it is to then build a system that one group eventually takes control or builds around. In some ways, the greater the less the governance, the less the risk of capture.

And the true value of DeFi therefore, can be found almost when the governance itself shrinks down nearly to the point of imperceptibility. Next slide.

Now, forgive me on that, what DeFi governance is not is something that can ever fully go away. And what we mean by that, of course, is there's always what's known as essential governance. Essential governance is the idea that there are certain things that simply have to occur, in part because there will always be human interaction with DeFi.

Until the AI machines replaces us entirely, there will never be a world where DeFi fully operates without any human involvement. There are still humans on either end accessing the protocol or encouraging smart contracts to access it.

The second is constant formality. Now, DeFi is in many ways thought of as token holders but there's also stakeholders that are different from token holders.

It is the case that not all stakeholders will

be token holders, but even though all token holders will be stakeholders. And these groups are always changing. The entities that use a DeFi protocol today may be very different than those used in the future.

That's already been the case for a number of protocols that have been nascent in the last few years. So our belief is the optimal approach to DeFi governance is minimization. The reason for this is in part, we think flexibility is critical.

These are game theory-based systems. They will -- the users will respond to how the system operates based on interaction and seeing how things change over time. So flexibility is paramount. Again, it's very important to distinguish between token holders, and stakeholders.

These groups will also change over time. In many ways flexibility is again for us the most important thing and there is a connection between a core mechanism and the human input.

There is consensus which is the idea of the Layer 1 itself, whether Ethereum, bitcoin or some

other network. They are the oracles or the truth validators. And then there's two things that may qualify as essential governance and maybe not.

One, which is treasury management. That's basically how the DeFi protocol will have a DAO control its treasury. Maybe that could be automated in a way, maybe it cannot. And then complex parameter setting.

And by that, I mean, things like choosing the collateral to use. MakerDAO, for instance, has used governance as a means of determining which collateral it accepts. Whether or not that can be replaced in time is uncertain. Next slide, please.

Now the thing I would also stress here is governance is not a panacea. There's a constant risk of hard forks. If you ever have a decision in place that a majority of users of a protocol dislike, they can always hard fork the system. That's innate to the idea of crypto. And changing governance somewhere in the system does not mean changing it everywhere. It is possible to change it in narrow fashions.

According to one of our co-founders, Fred Ehrsam, if the point of a blockchain is to provide a ledger of universally accepted truth, its integrity is paramount. In many ways, this is the core conceit of the blockchain.

When you start breaking down the integrity of it, the whole system falls apart. Next slide. Now the insights of regulation we think are four things. First, inessential governance is likely to be competed away over time.

We should know over time whether or not certain things are essential or non-essential. In particular, we're watching as I said, whether or not complex parameter setting and treasury management will qualify as essential governance over time. The third, data is paramount.

We have such little data overall how this all operates. There's a need for greater research. I know that's a common statement around here at TAC but it's especially true here.

Third, there is a distinction I think between principles, regulation and code. At the CFTC, the

classic regulatory regime has been principles-based but there are examples of using specific regulations. It is likely the case that the principles-based regime will most flex to the flexibility required for DeFi's governance structures.

And finally, I want to really stress the danger of what I call DiNo or decentralized in name only. There are a lot of DeFi protocols that claim to be decentralized but are not.

True decentralization is decentralization in reality and the goal of regulation in some ways should be to keep to that truth and try to reward those decentralized protocols that are actually decentralized while punishing those which are not.

With that, I'll turn it over to Ben for more thoughtful commentary.

MR. MILNE: I should have asked to go first. Well done.

So I think I'd come at this from slightly different perspectives. I felt like the best use of time was to try and bring us back to some first

principles in that a DAO is effectively an organization then hypothetically runs on smart contracts.

So what can you do with a smart contract? And the reality is, next slide, you can do a lot of the things that you can do in traditional systems.

You can just enforce those with smart contracts, alongside existing regulation and even existing corporate structures. Next slide. What I mean by that is in a traditional structure, you have governance policies and you store those somewhere.

That's no different if you're actually managing some level of voting through a smart contract infrastructure that uses a blockchain but the mechanisms under the blockchain and how the data is managed, is simply different. Next slide.

Key differences that you probably buy into if you want to be utilizing this new technology is that in a traditional system, we do things like this. We fly, we sit at a table; sometimes we vote. In this particular case, we're advising, we're not

voting, but if we do vote, many times it's still happening at a table. In a smart contract system, there is no table.

Everyone's signing in a globally distributed way. But that doesn't necessarily mean that who is signing or what they are signing is not regulated. It just means it is a digital first signature. Next slide.

Also, in a traditional system, I think we would all agree and like a board context many times votes are private by default and then you choose to share as a feature.

Many times in a smart contract infrastructure, vote is public by default and you choose certain things to be private, which is now technically more feasible than it was a couple years ago at least with blockchain infrastructure. Next slide.

So smart contracts can engage in regulated functions but not all do. My personal view is that if a smart contract engages in a regulated function, it should be treated as no different than any other computer code. Let's also recognize that

not all smart contracts exist and behave in regulated ways. Some things actually provide basic community guidance and management.

Next slide. Taking back to a traditional structure, let's imagine the board has a vote and maybe that vote is related to a bonus program. In certain companies, compensation is public. Next slide.

Utilizing a more sophisticated smart contract model to essentially gather and distribute the votes, as was mentioned before, token holders can be stakeholders. Token holders could be a board.

Token holders might also be a community selecting an initiative that has absolutely nothing to do with regulation generally speaking. Maybe they're voting on a community initiative that has to do with like whether or not to buy a golf course. It doesn't always need to be a regulated activity.

Now, smart contracts are interesting as was mentioned because they're also composable, which means you could compose one on top of the other.

And you can do so infinitely. Next slide.

You should know first that this code is gobbledygook. It's generated by a ChatGPT. It has no purpose. It's not audited here. It's simply for example purposes.

But if you are going to deploy a smart contract with addresses that are authorized to release a bonus, which was the example that I just mentioned, that's not an out-of-this-world activity for any board to approve and then utilize a mechanism to release the payment. But in this case, it's released in code.

That gets more complicated, next slide, when you start adding on the composability nature because a trade could be executed in real time the moment that actual decision is made available on-chain.

And so when you get down into the trail of things, it's very difficult to understand what the actual relationships are between these things and where it could originate with a smart contract that is not engaging in a regulated activity.

It could result in a smart contract's use of that information to then engage in a regulatory act -- regulated activity.

That is an interesting question to answer, but I think as we have just learned and hopefully many of us knew, if there is a person operating a piece of software, that person lives in a place where it has laws and jurisdictions applied to where that person is, or where the entity is.

So again my view is that as long as people are operating these, if they're engaged in a regulated activity, well, they are engaging in a regulated activity full stop. Next slide.

You might ask yourself with all this complexity, why in God's name would any technologist actually want to use these things?

Because it's actually much easier to take notes and get approval for the notes in the next board meeting, then to implement an audited smart contract, which could cost who knows how much money to actually implement basic votes.

So I put up a few things that if you buy into

them, you might actually buy into utilizing these technologies. And if you're not convinced, no amount of words would convince you. Next slide.

Something that's very unique in this moment in time is that the blockchain ecosystems and technologies are now at a level and rate where they're actually cost competitive with relational databases, which means in certain scenarios, they are now the best tool for the job, not just a cool new tool to utilize. Next slide.

Balancing centralization and decentralization for many companies or organizations becomes a balance of privacy and security, and it gets more complex than it was because new blockchain technologies can now be private by default, not just public by default. Next slide.

There's been a great deal of work done in private permission blockchains. I wish the team from Avalanche was here today. Avalanche has done a great job with private subnets where you could stand up specific nodes and manage those nodes in a private environment, where it's actually very

similar to a relational database with the right permissions wrapped around it.

In that particular case, it's just code to host smart contracts, which are composable with public blockchains, which is very unique. Next slide.

And so I'll put out a hypothetical, a really boring use of a smart contract in a current regulatory structure, where many participants could be known to one another and it's non provocative. Next slide.

In this particular case, there's simply the mechanism using a smart contract to gather votes by known stakeholders, known to everyone selecting that certain votes should be made public by default instantaneously which are composable for the public good. Now this has nothing to do with trading.

It has to do with utilizing smart contracts as a technology to get information to people who should have it faster. And I don't necessarily know anyone who's doing this. And it would be relatively expensive to do it, but I imagine because of the

benefits I mentioned at some point in time, someone will. Next slide.

Either way, the word I think was accountability was utilized earlier and in any of these -- I do hesitate to call them schemes -- but any of these designs, you have a signature. In a classic traditional system, you have wet signatures.

If I sign Elvis Presley on a page, everybody knows I'm not Elvis. But if I gave you that address, you might believe it belonged to me. It absolutely does not belong to me. However, you can programmatically verify who it belongs to with everything that they've ever signed.

And so, somewhere in these things are levels of accountability and proof of contribution which are quite difficult to prove in traditional systems. Next slide. Either way, there is a way to sign. One is just a different type of pen.

And so, for much of the smart contract infrastructure, there are hardware wallets that are carried around to sign and kind of a running joke

sometime is which pen do you want me to grab? Next slide.

Now, the immediate application of smart contracts are to, from my perspective, solve real-world problems probably less to do with how do we translate things like votes into a smart contract?

More like product creation, which to the circle team's credit has done a great job in productizing how you integrate with decentralized protocols, or even various versions of them to create new products that become smart contract-native. Next slide.

So when it comes to people, my view is at least that people live in specific jurisdictions and those jurisdictions have laws and there are nexus questions related to that that still apply regardless of what technologies you are using. Next slide.

And so there are a couple questions that I have coming out of this knowing that if people are going to be regulated and organizations are going to be regulated in the jurisdictions where they

operate, which certainly makes sense to me, what do we need to do in order to remove the amount of confusion over how the regulation maps?

From my perspective, one of the most important things we can do is agree on definitions. My personal perspective again is that one of the most important things that happened in the cryptocurrency industry is the definition of an exchange.

And then I think we have questions like how do we treat future regulation? When we imagine a world where a person is not involved in the creation of the code, it is entirely decentralized.

And there is no personal benefit, which is somewhat the next evolution of a self-creating DAO, hypothetically speaking, and maybe a fun conversation for later. Next slide.

If the presentation is useful, I've made it available for download on IPFS and a couple other gateways. I'll be around for questions. Thanks for the time.

MR. REDBORD: Justin, Ben, thank you so much

for a really terrific presentation. One sort of big takeaway for me I think a really great intro into this space is really two things.

One, you started with definitions and I think as the subcommittee sort of dives into its work here, that's going to be very, very important.

And the other piece is governance structures, right? Justin talked about spectrum, not a toggle, and then sort of dove even deeper into that. But I think it's so important to sort of understand all of those things.

And quite frankly, ask questions, like does it matter, right? Treasury recently said they might not from an AML perspective. So I think it's really -- it's a great jumping-off point, the definitions and sort of governance piece.

Would love to kind of open it up to the committee for questions or comments on the presentation and Nikos looks like he's going to kick it off for us.

MR. ANDRIKOGLANNOPOULOS: I think the interesting idea of being a spectrum is not just,

it's a spectrum, but it's also dynamic. So one day you can be a DAO and can be decentralized.

But if you're talking price kind of gets reduced, people might be start fleeing your organization or voting or not participating. So I think that brings up the question of surveillance and continued surveillance.

So here you have an organization that is far more dynamic than what we are used to. And I think when you think about surveillance, it's not only the legal formation. Who gives you the status of a DAO, but do you maintain it?

And that's kind of from a legal perspective but also from an operational perspective. Depending on what kind of services you are offering and you get your license to offer financial services, what are the requirements on that surveillance?

Making sure that in a decentralized world, you can offer business continuity. You can respect your consumers and protect the end consumer. So this is kind of my comment. I'd love your thoughts on how you think about surveillance and how that kind of

relates with DAOs.

MR. BIAGIOLI: It's a great question. I think the number one thing to think about is that any good DeFi protocol or system will be very susceptible to on-chain surveillance by everybody.

That is in many ways, part of the core value is that you should be able to see what's going on. Now, there's a -- it's worth noting a differentiation there because there is the actual protocol's usage on-chain.

And then there's the conversation about the protocol, which can occur in other places. And I don't think we really can gauge in how to deal with that.

For instance, a lot of DAOs use Discords or use Telegram or various other things for the nature of their communication.

And I think that's an interesting question is then that is often quasi-public because they're often available for anyone to access or is available for anyone to access who is already a member of the DAO, but they may not be as publicly

available as the chain itself.

MR. ANDRIKOGLANNOPOULOS: If I may add to this, sometimes, also you see those forums being influenced by certain members.

So once you are dealing with DAOs, there are certain members that can bring you around to the members that they consider most influential. So that basically you have a sense of whether the decision of the DAO will go through or not.

So real decentralization I think needs to agree on what are the metrics that maybe the regulator needs to be kind of monitoring or whoever issued the license and the formation of that DAO. So that it truly maintains the status of the DAO out there.

MR. MILNE: There's an interesting thing happening from my perspective is that it seems as though some groups have relieved themselves of understanding the regulation prior to starting the project.

And then you have collective voting without the awareness that they might actually be violating

a law. And so, in the event that there is an organization that starts something and then it accidentally or knowingly violates the law or tends to obfuscate it, there will probably be additional actions as those projects are discovered.

That being said, there are a lot of community projects based on smart contracts. Some might call themselves DAOs or something else that are just simply not engaged in regulated activity, at least that's my perspective.

And on the surveillance front, there are two sides. If you are engaging in regulated activity that it's like AML and BSA compliance, at least in the United States and I'm not sure where you would not be subject to it, but those are requirements of the organization offering the service.

And there are other ongoing risk monitoring services that requirements of a basic AML and BSA program that if you're going to offer those services, those things just have to be in place.

MR. REDBORD. Thank you so much. We're going to go first with Hilary and then with Gün, who is

online.

MS. ALLEN: Thank you. So I have a question for Justin about this concept of government minimization which in many ways is, right, a broader question about automation, isn't it?

I mean, so I wonder how much we actually want to minimize governance or how much we want to automate things because it works fine most of the time. But then, when it doesn't, then what are you going to do, right?

We've always needed sort of scope for intervention flexibility, et cetera, in unanticipated circumstances when things go wrong.

And so if we can't really minimize the need for governance in unexpected circumstances, because you can't program every eventuality into a smart contract, then essentially, you fall back on then the community to decide what to do in those circumstances, if it's truly decentralized. If it's not, then you actually have a decision maker and then it's centralized.

And I think about the fact that we've had the

technology for decentralized decision making for hundreds of years, right? A general partnership can be a decentralized body.

They have never really sort of grown to scale because it's hard to coordinate decision-making at scale and you get into conflicts and all kinds of disagreements et cetera. And so centralization has sort of been the natural convergence for scale reasons.

So I guess my question is, how does this work when you have tried to minimize governance but then something unexpected happens? Does it by default then have to become centralized when something goes wrong?

MR. SLAUGHTER: I think it's a good question because I think it depends in some ways upon the nature of the event. I would say we have examples of large decentralized social movements and that's the way a lot of political actions have taken. Your point about how often they persist over time is the most important one, right?

Is it possible to have a consistent

decentralized movement forever, and you can find examples of it in politics and policy, and things like political movements, environmental movements, rights movements. But it's very rare to have a formalized system.

What I would suggest is that it's possible that the solution is for government to focus on the actual people using it rather than the protocol, is one option.

In terms of what happens when things go wrong, which I think is the critical question, this is where there is, I think the greatest amount of research and engagement to be done, which is actually to look at how DAOs and how other things have responded to events going wrong and seeing how they moved forward.

Now, it's worth noting, right, it is very rare to see persistent decentralization, but in part that's because it's so easy absent consistent guard rails for things to fall into centralization.

Optimized decentralization is fundamentally a red queen race where credibility persists because

of consistent resistance to decentralizing forces. And I think in many ways, that's probably a role for regulation and government to the extent it's possible.

MR. REDBORD: Okay. Thank you so much. Gün?

MR. SIRER: Hi, everyone. And Ben, thank you for the shout-out. I regret I cannot be there in person due to family circumstances but very much enjoyed the conversation today.

Justin, I think you nailed it when you mentioned credible neutrality, that a DAO must be credibly neutral and I loved the DiNO word, the decentralized in name only. So my question to you is, how does one measure -- what's the metric by which we know when something is credibly neutral.

What are the kinds of analyses that one should set up? I'm curious about how you are thinking about them.

MR. SLAUGHTER: Not well, honestly. It's a complicated question. I think in many ways, it's probably a multiform analysis that is better left in some ways to policymakers than any individual

like me. That said, I think probably it's four things.

The first is probably the degree of dynamism in the network. If you have something that is functionally not doing very much almost by -- if it's a ghost town, you can question whether decentralization exists. The second is the diversity of stakeholders and token holders.

The third is likely the nature of how the system is maintained. Something that is actually managed by a small number of people with true formality could be questioned to be decentralized. And the fourth is probably overall the breadth of the system.

But these are very initial thoughts. I think in some ways the great question should be how do people in this room, broader policymakers define decentralization in terms of what are the goals we want to seek? And that's the best metric setup.

MR. REDBORD: It's so interesting. I think the question and the answer really gets to the heart of so much of the work we're doing, right?

I mean the focus has always been amongst regulators and policymakers on sort of how to regulate a centralized ecosystem. I think we're just really shifting to what arguably is the more interesting conversation is how do you regulate in a more decentralized space?

I think to some extent to that end we're going to go to our third presentation, Dan Awrey, Professor of Law at Cornell Law School, will present on the topic of stability and security challenges and regulatory implications for crypto, and Carole, I will let you also lead off with a question because I'm sorry I made you put your plaque down.

All right, Dan?

MR. AWREY: Thank you very much Ari. Hopefully you guys can see some slides. Despite the name -- the lofty name given to this presentation, it really carries on from the last session in many important respects.

And really what I want to spend my time talking about are how decentralized systems and

actors place stress on the design of existing regulatory frameworks to Ari's point a moment ago. Do you guys have some slides up?

MR. REDBORD: We do.

MR. AWREY: There we are. Excellent. So as all good presentations will, I will start with my own definition which is really a non-definition in many ways of decentralization. A few things to say about this to start.

One, none of this is to say that any particular element of decentralization is an unalloyed good or posts as a clear and present danger to the republic.

One of the big things that I think that we've got to grapple with is that most of these things have been around for a long time. In some circumstances, they have proven to be valuable. In other circumstances, they've proven to be problematic.

And what's really happening here with decentralization is actually a series of technological shocks that we've now on our second

meeting talking about in terms of automation generally, in particular the automation of state-contingent contracts and smart contracting and then increasingly AI.

But the way I think about decentralization is one along five dimensions, the first being development. So who is making the thing? In the case of a lot of the DeFi space, who is writing the code? Who is designing the product and who is setting that down for the purposes of presenting it to the world?

Now none of this is particularly new, right? Even generally you can look to open source software; you can look to open APIs and even in finance there's elements of things like open banking, which have since their advent, been in effect a decentralized development space.

The second dimension is governance. We've talked about this quite a bit already. Here, there's really two axes. One is how much is left to the residual discretion of human beings?

In a very automated world, there may be very

little but there may be still decisions that remain with human beings and understanding both the numerical balance, but also the importance of the decisions that are automated versus being made by humans is important.

And that governance can be one human with a button; that governance can be dispersed amongst many different actors and stakeholders. It's just another dimension of decentralization and one that itself has a lot of different permutation.

Three, one of the things that I think is more interesting here where automation in particular comes into play is operational decentralization.

And I'll include in this category as well transactional decentralization, which is really just one specific form of operational decentralization. Traditionally, I go to my bank, whatever I want my bank to do, the unitary actor in that system is a bank.

And insofar as they engage in relationships with third parties or outsource certain elements of that relationship, the absence of contractual

privity means that, whatever that relationship is it is one governed by regulation and I don't particularly care about it as the end consumer.

And as long as we're talking about one product or one service that can also work fairly, but one of the things I'd like to talk about in a moment is when we get into a world where there may be a single activity but multiple different actors who are providing inputs to the activities, so the finance supply chain if you will.

How it is we allocate regulatory obligations and overall responsibility within that supply chain. Fourth, we have balance sheet decentralization, right?

This gets into Justin's earlier observation about custody where really, since the dawn of time, time always starting with the old English bailment case law if you're somebody like me, we've really been dealing with decentralized actors.

People who we can identify were responsible for the custody, the holding of a particular asset or in exchange for one asset issued another asset,

some sort of promise to the customer. Once we get into a decentralized balance sheet world that may not be the case.

The promises may come from multiple actors. The party that is holding the asset may not ultimately be responsible for delivering it back to the customer. And then we have things like self custody where fundamentally we have to look to new legal relationships in order to understand how that works.

All of which is to say that our definitions here are not particularly tractable but they're very important in terms of understanding what we think is important, which of these new developments are actually new, which of these new developments are important?

And the one that I want to focus on here really is the fact that across all five of those dimensions I just mentioned, existing regulatory frameworks have relied historically on a very, very high degree of centralization of the actors and activities involved in financial services.

So we can look at organizations, regulators like the CFTC, sister institutions as engaged in a series of tasks, right?

So identifying outcomes in accordance with their contractual mandates, writing and updating rules that are designed to achieve those outcomes, monitoring compliance with those rules and then enforcing against any breaches of those rules.

And historically, if we can go to the next slide, an important and often forgotten feature of this system is that organizations like the CFTC actually delegate an enormous amount of the heavy lifting to regulated actors themselves.

So you can think of a couple of examples here. You can look at risk based rules under the AML CTF regime. You can look at risk management rules for derivatives clearing organizations. You can look at the development and execution of things like custody rules for financial assets.

All where the regulatory burden is split between the regulator and the regulated institution. And here centralization is great on a

number of levels because it makes it easy to do two things.

One is for the regulator to say you, yes, you have to go out and find ways of implementing these rules to our satisfaction. Two, if you don't, you are the person that we are going to hold responsible for any failure to actually live up to our regulatory expectation.

But what this means, that delegation of responsibility, is that we have a particular ecosystem when it comes to regulatory compliance and financial services. The best example of which is probably AML CTF rules.

If we didn't delegate at least some of the design and execution of these rules to regulated financial institutions, FINCEN would be the primary employer in Washington, D.C.

The number of human beings in the United States involved in regulatory compliance for major and minor financial institutions is enormous and dwarfs the federal government as a whole.

What the implication of that then is if we

can't find ways to find a constructive way of delegating some of this to financial institutions in the decentralized space, we're going to have to fundamentally rethink the way that we approach regulation. Next slide, please.

Once we get into the decentralized space now, then -- we've got to then grapple with the fact that it becomes more complicated to allocate this responsibility.

Once I dice up components of the supply chain, it makes it more difficult for all the elements in that supply chain to share the same information, or I have to think about ways of designing my ecosystems in a way that made sure that information is available to all in a timely way.

It makes it more difficult to identify exactly who is responsible for any given regulatory failing. That is to say, if something goes wrong in my complex system, it may not be immediately apparent where the failure actually lie.

And when things do go wrong in automated systems, decentralization means it might not always

be clear, or at the very least, we have to think about it in advance who's responsible for intervening to effectively take action to flip the kill switch and decide how to react to a problem that was not specified in the software that up until that point had been running smoothly.

All of these things, if we could go to the next slide, then leave us with a series of questions that we've got to think about as we think about -- not just about what DeFi is, but also what regulation needs to be to make sure that DeFi is something that is constructive, something that does not create opportunities for military arbitrage. Something where the consistency, in terms of regulatory treatment and outcomes for consumers, is equivalent to TradFi.

Here, I often talk about what I like to call the first law of regulation. We hear a lot these days about the kind of glib statement, same function, same risks, same rules. I actually think that's a really poor guide for regulatory action.

Not because I don't think that we should be

thinking about how different types of financial markets and institutions perform similar functions, or that we shouldn't be thinking about how regulation needs to evolve in response to that, but because at the end, where it says, the same rules, we automatically put our previous experience and the tools that we currently use to use in ways that may not be particularly helpful.

An example I'll give is talking about deposit insurance for stablecoin. Right? Once you decide that they're functionally equivalent to a bank deposit, people automatically jump to the current regulatory framework for understanding how we regulate bank deposits. But in reality, there's lots of different strategies, some of which may be better suited to regulating stablecoins.

And what the first law of regulation does is to say, look, what we're trying to achieve stays the same. But the way we do it, we're not necessarily going to use the same rules. In this respect, it is like the first law of thermodynamics, energy can neither be created nor

destroyed, it can just change shape, change form.

And really here, this is what regulators are trying to do as well to understand how the evolution of finance in response to technological shocks actually requires us to rethink some centuries old path dependence in the way that our regulatory systems are designed.

So listed here are what I think are the five key questions. First, how do we make the regulatory perimeter make sense in this new world? This is something with the Ooki DAO that we talked about a moment ago that we had a person for the purposes of an enforcement action.

The bigger problem though is what happens if I have a decentralized actor that only does one part of a regulated activity? Is that going to be sufficient to bring them within the perimeter of regulation ex ante?

Or am I going to be perennially forced as a regulator to use enforcement action as a substitute for ex ante regulation there?

As we look across the CFTC and other

regulators, the answer to this question is going to change a lot. We have very different designs of regulatory perimeters in the United States.

The definition of a bank is very different than the definition of a systemically important financial institution for FSOC designation purposes, which is very different than the way that we look at something like the definition of a security under the Securities Act or Securities Exchange Act.

Each of these perimeters threshold legal questions, because they're designed differently, are going to have different answers in terms of how robust they are to the emergence of decentralized actors and activities.

Second, once you're in the regime, how do we allocate responsibility for taking action when action's required outside of the automated ecosystem? And secondly, allocating responsibility for when things go wrong. Who pays the bill in the event of regulatory and market failures?

Again, there's going to be no single actor

here. But what this challenge requires is looking at these different ecosystems and try to come up with ways that one, clearly sort of indicate who's going to be responsible so that people, including the decentralized actors, can govern their activities accordingly.

And then two, trying to articulate the answer to this question in a way that is robust to changing circumstances. If we should've all learned anything over the last several years, it's that if you try to regulate an ecosystem at a particular moment in time, you're ultimately probably going to find that the regulation moves much, much slower than the ecosystem it's trying to regulate.

And this leads to my last two questions, ultimately, which I think are huge questions, one that we've struggled with since the creation of the abacus, and really haven't spent enough time dealing with in a fundamental way, especially in a system like the United States where we have a fragmented regulatory community that is often slow to act, often for good reasons, in response to new

technological change.

So first, how do we get regulation into automated systems? To some, this often seems like a trivial question. Well, I just code it, right, do X, don't do Y, into the system. And that presupposes a very particular type of regulatory rule, mainly, a hard and fast bright-line rule.

But as was mentioned earlier, a lot of CFTC regulation and a lot of financial regulation in general is actually either outcomes based, in the sense that it's designed to be open textures and require analysis of changing information and circumstances. Or it's simply designed to lead to a particular type of behavior that is then left unspecified, and potentially open to ex post adjudication.

But if we want this to work, regulators, technologists, firms need to understand what type of rules can be embedded in code, where we're going to have to put the human discretion in these systems, and how we're going to deal with the fact that rules often change. And that when rules

change, somebody will often have to go back and take a look at what needs to be updated, and how to integrate it into these systems.

And to finish off, just to give one example here, it's often put to me that decentralized actors can embed oracles within their system so when the federal register updates regulation, that can just be inputted into the existing regulation. Basically, a software upgrade, in effect.

But again, what does that mean in a world where it wasn't -- the rule change was not fully parametrized in a way that is possible of articulating in code? What happens if it requires you to think of possible solutions to problems as a precondition to implementing them?

What if, for example, to use something that was discussed earlier, you're dealing with situations where you want to maintain credible commitment to always be in compliance across very different -- various different outcomes, including things like forks where one of the nightmare scenarios we might envision is a fork where the

fork is based on one fork wanting to comply with relevant law, and the other fork deciding that that wasn't for it. Right?

Those sorts of things are all things that, while some are trivial, some are more challenging, all need to be sorted out strategically in order to make sure that you can comply with the first law of regulation.

And this doesn't go to the level of prescribing the individual rules, but it goes to having a sensitivity to the fact that whereas previously, we were in a situation where it was very easy to call up the chief compliance officer at Institution X, and talk to them about how it is they would implement the regulatory -- the regulative principles and rules for their particular organization.

How much of that apparatus can we, should we keep in the context of decentralized spaces across different actors and activities? So I don't propose to put any answers on the table today, but one of the things that I'm really looking forward to on

the TAC, and especially as we look at the subcommittee on blockchain and digital assets, is to explore some of these questions in greater detail.

And to have the technologists, the lawyers, the regulators, and others sort of bringing their probably quite different perspectives on the importance of these questions, and what they think the best path is forward. So thank you very much for you time. I look forward to the discussion on this.

MR. REDBORD: Dan, thank you so much for a great presentation. And really, that thoughtful ending around sort of Justin starting with definitions, and then Ben into sort of governance, and now sort of what regulation could potentially look like, or at least the questions we should be asking in a more decentralized space

I think that's -- that's what's really exciting about the work of the subcommittee. We're going to open it up now for a few minutes for discussion. Carole, do you want to kick us off?

MS. HOUSE: Sure. And honestly, after the question that Hilary asked, which also mirrored mine about governance minimization and, like, that being an optimal approach. And let's define optimal, because there's goodness and resilience that comes with decentralization, but then also the issue of accountability being a real challenge, which then Dan got into.

So ultimately, I'll just say that I -- how much I appreciate the varying sort of perspectives and issues that each of the three presenters spoke to. And, Dan, how much I appreciate talking about the law of conservation of regulation.

I know I've mentioned it another context too, sometimes that there would be folks in a space not necessarily represented here who felt that, well, if we can just decentralize enough, there will be no responsibility pushed anywhere. And ultimately, I don't think that's a reality if the risk landscape moves.

Regulatory frameworks are risk-based, just like AML programs are, our anti-money laundering

programs, and other, like, cybersecurity programs for banks and such are risk-based approaches. So if the risk landscape adjusts, I would expect regulators to do so.

One observation, and, Dan, I appreciated you going through all the different functions, like you mentioning governance, and operations, and transaction development, etc. I know that that triggers a lot of really interesting questions for folks about, like, where -- then where do you want accountability to sit?

Developers certainly get very uncomfortable when they think about where -- about accountability sitting with them for developing code. I -- but ultimately, if you're not relying on the application layer, which is where most regulation ends up sitting, and the argument for many folks to focus on these centralized -- for regulators to focus on centralized exchanges.

The, like, three areas that I see that are ripe for choice, I may have mentioned in the first TAC, are network, protocol, and user layer, the

people responsible for some of the development, governance, operations, etc., and what does that look like?

What does, like, pushing responsibility or accountability to them look like? Does it end up sitting just with end users, and that they're the ones responsible for understanding what they're interacting with? What are the implications then for developing DeFi systems that -- where consumers can actually understand the decisions that are being made in the system that they're interacting with?

So just more of an observation. Not a question anymore, since I think your presentation really hit on and outlined a lot of the key questions that I see. But I'll turn it over to, I think Nikos.

MR. ANDRIKOGIANNOPOULOS: Yeah. I want to make -- I want to make a comment on the state of how I -- the world this day on decentralization.

And one example is, when we think about Ethereum ETFs, and we think about financial institutions performing custody on Ethereum,

performing trading, and performing all of those functions, there was a recent event, like, a month ago, where Ethereum could not finalize transactions for half an hour.

And during that time, in a normal world, if that would happen at the New York Stock Exchange, there would be alerts, there would be an entity, there would be calls, there would be coordination. There would be a bunch of mechanisms that would be triggered to protect people transacting over the exchange.

In the Ethereum world, that doesn't exist. So in the financial institutions that are going to be providing those services and will be applying for crypto licenses, I think there needs to be embedded monitoring and surveillance there to make sure that they protect the consumers.

And to Carole's question, I think we need, I think, to embrace the fact that a lot of the responsibility lies at the application and the service providers that need to protect the consumers. The continuous nature and the

accumulated risks on the decentralized entities, there need to be protections on both sides. And the question is, how much on each side?

But I think with the -- today, as we embrace those things, we are still thinking in a centralized world, and I think we're missing on how these things work. And if it weren't for half an hour, and it was for 17 hours, what would have happened then?

I think these are the things that's the reality, I think, that we're facing today as we walk into that.

MR. REDBORD: Dan, or anybody else? Any comments on that?

MR. SIRER: I might -- I would love to follow up on Nikos. And -- if I may.

MR. REDBORD: Sure. Go ahead.

MR. SIRER: Okay. Yeah. Nikos, that's a-- that's a very insightful question and a very interesting scenario. And I think the -- one of the big pitfalls I see in this space is to see the L1s, or is to see the blockchain as a financial entity,

a singular entity.

And I think the right way to view them is more akin to how we view the Internet. When the Internet is down, and we have many Internet outages per year all the time, we don't hold the Internet responsible for the quality of service, we don't hold -- we don't expect the Internet to uphold a packet to arrive to its destination.

We rely on the endpoints to make sure that there are multiple redundant paths to build multiple redundant paths into their backend operations and so on. And, of course, to build the kind of monitoring that you mentioned into their daily operations so that outages in the Internet don't affect the end users.

And that is the thing that gave us the Internet explosion, that end to end design. And blockchains are very similar. They carry out functions to the best of their ability. There are multiple L1s, there are multiple systems that can provide approximately comparable services. And it should be up to the endpoints to provide the

service. And the moment we see -- start seeing things this way, I think we have tremendous opportunity to innovate, because the networking side becomes an open fabric, a common space that carries out operations when it can, and it's up to the endpoint, the edge points to uphold the -- whatever the service level requirements are.

MR. REDBORD: Thank you so much. Dan, were you -- did you have a comment also, or a response to Nikos' comment?

DR. AWREY: No. I think I wanted to voice my agreement with Nikos that I think mapping how we think about the system is an important first step here. And then understanding -- and part of it is a priority that I have that in any complex system, failure is inevitable.

We just -- especially human systems, which these ultimately are. We have to acknowledge that things are going to happen. And then the question becomes, where are the important threat vectors? Where are the important failure points in these systems? And then what sort of things that we might

even take for granted in the context of more centralized systems do we need to consciously build into these centralized systems?

And I think Nikos raises a good example of this, where the New York Stock Exchange is going to let me know quite quickly, whereas I'm going to learn about problems on the Ethereum network via social media networks in a sort of circuitous path that is going to be individual to each particular user.

There are opportunities in that, in the sense that there are, as Gün, I think, is referring to, ways to rationalize the system and make it work better, or these systems and make them work better. But we shouldn't be any less cognizant to the fact that we're trying to build pretty big, complex things.

And in finance sometimes, unlike the Internet, the ability to have those things when we need them, and the difference between my ability to execute a go- -- Google search and pay my rent is big and important, and may mean that as we start to map out

what these systems look like and where their vulnerabilities reside, that it's not just the network design that's important here from a technological spectrum, but also what we're doing.

And finance involves big, important things where oftentimes there's no take-backsies. And in that world, it may be that we want to introduce redundancies, security features, stability features that have the look and feel of what we're used to with more centralized actors.

MR. REDBORD: Thank you so much. I'm going to turn things over to Carole for a comment, and then you can take us into the next session.

MS. HOUSE: Sure. Thanks. I wanted to comment, because actually something that Emin mentioned really harkened back to an issue that Ben presented on, and I think Justin really hinted at it too.

The fact that these systems, and smart contracts, and decentralized networks can engage in non-regulated activity. And I mean, like, not just financial unregulated activity, just not financial activity. They can transfer information. Like,

almost all these networks can conduct an information transaction at the same time that they can conduct a financial transaction.

So ultimately, that issue needs to be something that's accounted for in, like, some of the -- whether it's embedded, or coming up with a perimeter, or the things that are built into, the policy, frameworks, governance, etc., to try to determine when there might need to be certain protections.

And why that -- why Emin's comment pointed it out to me, I felt, is that he's right that Internet service providers, like, don't fall underneath our sanctions frameworks in the sense that, like, them providing Internet access to a sanctioned party isn't something that would be considered a prohibited transaction.

The issue though there is that, like, ISPs aren't viewed the same way as, like, operating a payment network. And so this issue where these networks can facilitate both regulated financial activity that may include prohibited transactions,

like, under sanctions, that's a term for the sanctions regime, but then at the same time, be able to support information transfer activity.

That's really interesting. And it's a challenging policy and technical issue that I think will be -- will be interesting. So if you guys have any reaction to that, happy to take that. Or we can just go on into our third awesome topic about cybersecurity.

MR. SLAUGHTER: I actually think --

MR. AWREY: Can I -- can I just --?

MR. SLAUGHTER: -- think that's a really critical point. You first, Dan. Sorry.

MR. AWREY: Oh, sorry. I was just going to say, this is a good point on this, and this is why I kind of beat people over the head with the law of conservation of regulation, is that what's the non-Internet Internet, in effect. Right? Libraries.

Libraries presents some risks, but they don't really present that many risks relative to the Internet. I think we could all probably agree with that. Once we get into finance, we have to deal

with the fact that there is this parallel system running on a different network design that we already regulate to do certain things.

And on the assumption -- now, we can certainly do it better. And we should always be striving to regulate the centralized system better. But to adopt a radically different end user is the ultimate bearer of all responsibility -- the system here has enormous consequences that don't exist for the Internet, because libraries and the Internet are not great substitutes for one another.

And you'd be pushing a lot of risk into a system that was that fundamentally different, in a way that I think we need to talk about as a committee. Right? What are the points of destabilization?

Are we really talking about, if we have one system where we have rules that are designed to allocate responsibility to a set of professionals, to a set of regulated actors, to sets of people that we can relatively easy see, talk to, and sue, relative to a situation where we are going to

really inverse that burden, and not put it on regulators, not put it on regulated institutions or actors, but put it on the end consumers of these products?

That's a very different world. And to have the two worlds operating simultaneously is something that I think warrants a lot more discussion than we've given it to this point.

MR. SLAUGHTER: I actually -- that's a really good point. I would draw a distinction, though, between the Internet to library world, and digital finance to traditional finance.

I did something yesterday, which probably I'll get my head of security really mad at me for, which is I was on a street corner, and I paid for a airplane ticket by listing out the numbers of my credit card. If someone was listening to me, theoretically, they could have taken up on it.

But there was no risk on that, because of course, it was ephemeral. No one was recording, no one could see it. We all know that under the Internet, you have to be incredibly careful with

your own data.

And we've put a lot of risk on people, basically, where if you go out and say your real name, your home address, your credit card information, that can be massively dangerous very quickly.

To some extent on the issue of the ISPs, that was a policy choice we made in the '90s, where we were going to allow them to have reduced liability risk for violating sanctions for various things, because we regarded it as impinging.

But I think Dan gets at a key point though, which we have to consider what the risks are and what the various issues are. A lot of how we will consider the risk on end users depends on how many end users we think, and who the end users are will access DeFi themselves, versus will use the gating functions of CeFi exchanges.

And that is a real question, I think, for the polity in terms of what that might entail. I mean, personally, I kind of take the position that if you're able to do self custody, and you know, you

are willing to ta- -- with the wherewithal to get involved with that, I think it's okay to access DeFi with certain minimal requirements.

But I definitely think it's probably the case that if it gets too easy, maybe that would be a problem. These are hard questions.

MR. GUIDO: I agree. They're extremely difficult questions. And I'm a lazy person. So I always think, like, how can I get what's already in place to do the work for me? And it feels as though on one side, if these systems do go down for half an hour, it seems really terrible.

But a lot of the regs we have in place in the United States, like Reg E, Reg Z, so on and so forth, have already accommodated for the profile of the provider with downline failures. It's not uncommon for, like, ACH files to be held up overnight.

It's not uncommon for a downline provider to be failed because their Internet connection is not working. So my personal perspective is some of the existing regs and rules, if we think about that

mapping, they may already account for a significant portion of the potential technical failures.

MS. HOUSE: Thank you so much for a great discussion. And thank you, Ari, for turning it over. So we are now ready to explore our third and final topic of the day, cyber resilience for financial markets.

To begin the discussion, our first presenter will be Kevin Greenfield, deputy comptroller for operational risk policy at the Office of the Comptroller of the Currency. He is presenting on third-party relationships and interagency guidance on risk management. Kevin, over to you.

MR. GREENFIELD: All right. Thank you. If we could advance the slide. So just a little background. Third-party risk management has been a key focus of bank supervision, prudential supervision.

And when looking at it from a cybersecurity, cyber resilience, the level of dependency that financial institutions have on their third parties for cybersecurity and for resilience is important.

So I'm going to deviate from the presentation a little bit in giving some more references towards cybersecurity and resilience. But really, this is overall third-party risk management guidance that addresses all risks financial, consumer compliance, as well as operational risks.

This is not new guidance. This is actually based on the OCC's 2013 guidance on third-party relationships, risk management guidance. We work together with the FDIC and Federal Reserve to harmonize and put out a single interagency third-party risk management, updating what we had had with -- for modernizing the language.

As well as addressing -- over time, we put out several frequently asked questions, because every third-party relationship is unique and there is always going to be nuances. And we tried to address those, and we tried to incorporate some of the common themes from those FAQs in this guidance. If we can advance the slide.

And really the important thing with any third-party relationship, it is risk-based. So what the

guidance communicates is sound risk management practices. And one of the things that you'll hear from all the regulators on this topic is that risk management and control structures need to be commensurate with the risk of the activity.

So depending on the size, the complexity, and the risk profile of a given institution, what are the services? Clearly, we're going to look at payment services much more differently than the landscaping contract that the bank has for their branches.

But really understanding, and then what is the nascent -- nature of the relationship? Because there are well-defined outsourcing contractual relationships. But there also are a lot of business agreements and partnerships that can result in risk to the bank, and could put the bank in jeopardy. The guidance is also very much focused on what are the roles of executive management in the board? What are the characteristics of critical activities?

Because very much looking at those critical

activities, and then looking at oversight of the relationships that support high risk activities, including those critical activities.

And then most importantly, this is not a one-time-and-done activity. Whenever you engage with a third-party, it really is a risk management lifecycle. So we can go to the next slide.

No modern guidance is complete without a graphic, so we added one here. But this one is very important. And when we talk about the guidance, we talk about this lifecycle. And I'll go into a little more depth of each area, and what are some of the key characteristics, as well as then applying it from a cyber or an operational resilience perspective.

But when looking at this, I really have always referred to this as the five Ws. Whether, and that is when looking at the planning and risk assessment, whether this is an activity that should be outsourced, whether this is a partnership that -- or business relationship that the bank should engage in. Really those first steps.

Who, due diligence. Once that decision's made to go forward, deciding who the best partner is, who the best firm is, doing your homework, doing your side. But really determining who is it going to be with.

The next W is what. The contract is essential. I'll talk a bit more about -- contract doesn't guarantee anything is going to happen in and of itself. But if it's not in the contract, and outlined in that contract, you have no expectation that it will happen.

So this is very important when looking at cybersecurity controls, ongoing monitoring and testing from a resilience standpoint. If you're not outlining that in your contract, then following up and having that expectation of your third party is going to be very difficult to enforce.

The next W is the watch. And really this is, again, just because it's in the contract, it doesn't guarantee it's going to happen. What is your ongoing monitoring? How do you assess compliance with service level agreements? How do

you assess the controls and risk management frameworks that your third-party has in place? And what are some of the mechanisms that allow you to do this?

And then my favorite W, the why. And this is also very important. All good things come to an end. All third-party relationships will essentially terminate, and this is the termination phase.

And this is the one where there can be a lot of difficulties encountered, because it's not something that's often considered upfront when designing the contract of what does termination look like? What are the duties and responsibilities of each party?

What support is available to move on to another vendor or to another third party? All this is something that really needs to be negotiated on the front end, because I can tell you, it can be very painful trying to negotiate that on the back end when it's not defined.

So going on to the next slide. Digging a little more deeply into this, and we talk about

this guidance, planning really is the risk assessment. And this is really where banks need to make the determination, what are they willing to allow to operate outside the four walls of the bank? What data are you going to release?

What responsibilities are you going to put on that third party for cybersecurity, for operational resilience? And is that within the bank and the board of dire- -- the bank management and the board of directors risk appetite?

So this is something that really needs to be thoroughly considered, because a community bank outsourcing its wire activities, yeah, many do that. A globally systemic bank outsourcing an activity that is their U.S. dollar payment and clearing systems, there's going to be a lot of thought put into that, and then what kind of controls would you expect for that? So really looking and making those decisions.

Due diligence, I go back to really doing your homework. And again, this is commensurate with the complexity, size and complexity of the bank's

activities, as well as what is the actual service being outsourced or being engaged in through a partnership.

Again, this goes back to risk tailoring and this assessment being commensurate with the risk. We look at it as a community bank. Doing this assessment, we want to see them do the fundamentals, we want to see them safeguard the bank.

Again, when you see a globally systemic bank, that if that operation does not op- -- does not process as intended, there can be ramifications in the markets, and can -- in -- not just affect the bank, but in fact, affect the sector, we're going to expect that bank to do much more thorough due diligence for those critical activities. So it is very much a risk-based approach to this guidance.

And then contract negotiation. This is an area we were very detailed. And the important thing here is this guidance is not a checklist, it's not a rule, it's not something that we are -- require the banks to follow every single step. But in this

area, we put in a lot more detail, because, again, if it's not in the contract, there's not an expectation that it will occur.

So banks, and I'm sorry, I will always say banks as a banking regulator, but any financial institution or organization, really do need to be thorough in their contract negotiation, and really look at those terms, and make sure they have everything covered.

And this can range from treatment of data or information you provide, is that something the vendor can -- or this outsourced com- -- or third-party firm can use? Is that something that they can use for their own analysis and marketing?

Or are there strict rules around how that data is used, how is it secured, and what are their responsibilities if there is a pote- -- if there is a cyber breach? Making sure those things are defined are going to be very, very important.

And again, subcontracting, we talk about third-party risk, and I'm sure everyone's familiar, fourth-party, fifth-party on what are some of the

requirements, again, that you've put in place? That has to be in the contract. If we can flip the slide.

And then the ongoing monitoring part goes back to -- I cannot tell you how many times I've been in a room where there has been an operational issue, and I hear everyone saying, but that was in the contract, they were supposed to do it.

Just because it's in the contract, there's no guarantee. And this is really where we set expectations for ongoing monitoring, use of independent audit reports, SSAE teams, on-site visits, monitoring and reporting on service level agreements, regular testing of operational resilience plans, penetration testing, and reporting out on the results of that penetration testing.

Really, those are the tools a bank can use for ongoing monitoring. Because there are limits. There's -- the value of outsourcing or engaging with a third party go down if you have to build so many controls and so many mechanisms that the costs

outweigh

So it's definitely a risk-based approach. But depending on the risk of that activity, you're going to want to have those controls that you're satisfied the third party is conducting the operation as contracted and as intended.

And then with the termination, this is something that really want to again emphasize is when ending a contract, what happens with the data? What is the level of support? What is the cost associated with moving from one --?

I can tell you, I've seen many instances where mergers and acquisition activities were ceased because the costs of moving or terminating a contract with a third party became so large that it was no longer economically viable to continue with that.

And at the time, none of that had been defined. And that third party looked at it as, this is my opportunity to renegotiate and get as much revenue as I can out of this relationship, especially if it's going to be walking away. So

having those considerations in that built in on the front end is really important.

So go on to the final slide. So what are some of the key takeaways? And first and foremost, it is ultimately the bank or the organizations, it's ultimately their service, their product, their customers.

I can tell you, I have been involved from a bank supervision side or third-party supervision under the Bank Service Company Act, I've been involved with that process for 25 years.

And I can tell you, I have never, ever seen a third party bre- -- a third party be breached for bank information. I have never seen a third party failed to execute or operate on a bank transaction.

I can tell you, I have seen multiple examples where the bank was breached, or the bank failed to execute a financial transaction. And their third party helped them get there, and they paid a fee along the way for it.

So it's -- always when looking at these activities, it's ultimately your organization's

product, service, customer, you just cannot throw it over the fence.

And one of the guiding principles we always have when looking at third-party risk management guidance is you can outsource the activity, you cannot outsource the responsibility. And that's really important when financial institutions are looking at outsourcing or partnering with third-party firms.

I can also tell you multiple examples, because looking at it from a cyber perspective and an operational perspective, but a number of times where banks have run afoul of consumer compliance laws, or unfair deceptive activities. And that was a result of what their third party was doing.

But at the end of the day, the third party was doing it on their behalf, whether with or without their knowledge. And ultimately, that financial institution is responsible for that. So that's very important.

The other thing is, this was not vendor management guidance. This was not outsourcing

guidance. Third-party relationships can take many forms.

It doesn't need to be an outsourced relationship, doesn't necessarily even need to have a contract. But where you're engaged with a third party, and that can present risks to the financial institution, it's important to manage that and follow that lifecycle.

So with that, I went through that quickly, because I saw that we are running a little bit behind. But I'm helping with getting us on.

But really overall, I'll tell you that the presentation itself is pretty boring; I know, I wrote it. But really very much want to engage in questions that you all may have.

MS. HOUSE: Absolutely. Thank you so much, Kevin. And we're going to go through our presentations, and then move into a consolidated discussion session.

For our second presentation on cyber resilience, we have Hilary Allen, professor of law at American University's Washington College of Law,

and Tim Gallagher, chief security officer at Nardello and Company, jointly presenting on the challenges with understanding cybersecurity risk and implications for operational risk regulation.

MR. GALLAGHER: Great. Thank you, Carole. Go to the next slide. Good afternoon, everyone. As you heard, I'm Tim Gallagher. Since the last time we've met, I've changed firms. As you heard, I'm now with Nardello where I'm serving as the chief security officer and a managing director in the cyber practice.

So as I said before I -- in addition to advising my clients on cyber risk and operational risk, I'm also working to mitigate -- identify and mitigate our internal risk. I'll pre- -- presenting with my partner Hilary Allen here today on the threat environment and the regulatory landscape regarding cyber security matters. Next slide.

Quick scene-setter here. My background, as I've mentioned in this forum before is with the FBI where I worked financial and cyber crimes for over two decades. When I first came into law

enforcement, working cybercrime matters entailed using law enforcement tools, whether that be subpoenas, FISA intercepts, search warrants to identify the whole conspiracy and take people into custody.

Get them off the streets. Disrupt the organization and dismantle them. Seize their assets. As we know, as time goes on that we can't put our hands on a lot of these criminals -- a lot of these threat actors.

So as -- as my career progressed it was about prevention. Getting the information. Getting intelligence. Getting it out. Pushing it out it out to the financial sector. Pushing it out to the general public in the form of public -- public service announcements saying, hey, here's what you need to do to lock down your system.

Here's how you can prevent -- here's what you can do to prevent yourself from becoming a victim in the first place. Obviously, that worked to a certain extent. Then it became pushing out signatures, indicators of compromise, which we're

seeing right now.

Getting them out to -- getting them out to the private sector so that you can -- you can harden your system and keep the threat actors out. Fast forward now, advising Fortune 500 companies in my current position. You'll see the threat environment out there that we're dealing with.

Cybercrime was responsible for over $10 billion in losses in 2022. Having overseen the IC3, when I was in the FBI cyber division, I could tell you that figure is probably significantly understated. We could all say that with -- as my former FINCEN colleague there is nodding away, it's probably, a multiple of that. Yes.

Correct. $300 million fraudulent sign on attempts in the cloud every day, according to Microsoft. 53 percent of businesses have experienced a third-party breach. Kevin Greenfield just talked about third-party breaches, right? So there you go; 53 percent.

In that last slide there's a typo in it. It says 60 percent of businesses. It should be 60

percent of small businesses that experience a cyber attack close their doors within six months.

So that's what we're up against. The first three bullet points show how creative, how relentless, and to a certain extent how successful threat actors are in this cyber realm. The last one shows us what's at stake.

Go to the next slide. So where does that bring us now? I said from enforcement to prevention to resilience. You know, that's what it has to be about -- about resilience right now.

This is where the discussion will go back and forth between best practices and maybe what's required of you by your regulator. But from where I sit there's certain minimum standards that we need to incentivize resilience without being overly prescriptive.

Not a one size fits all approach. But has flexibility built into it so it matches your business model to your activity. As well as being appropriate for the sensitivity of the information that you're taking in and maintaining and handling

on a daily basis.

So what is resilience? From where I see it, it's the ability to prevent, withstand, and recover from a cyber-attack. And three things that I recommend to my clients on an engagement-by-engagement basis.

Obviously, there's many as far as resilience goes but here's what I put out there. Because as we all talked about, it's not a question of if, it's when. And it's your ability to bounce back once you've been hit.

Preparation. This all seems basic to probably just about everybody in this room, but it's not. Preparing. Having a plan in place. Most of the companies out there that we deal with don't have a plan in place.

It's pretty basic. Have a plan. Test the plan. Try and break the plan. Make amendments to that plan. And have it so that everyone in your firm knows what it is. They know what their role is within that plan.

Replication. Backing up your data. It's not

just backing up your data and putting it in a warehouse somewhere out in Kansas City. No. It's actually having it that you can actually stand your system back up when you get hit.

It's testing that data, making sure it's not corrupted, make sure you that you can put -- get your data backup, get your system back up online in a timely manner because if it takes too long for you to finally get it back up and running, you may not have a business anymore.

And then recovery. These are the tactical steps that you need to take, step by step to -- for the -- to get your short-term business plan back up and running again so that you can -- you could actually do business. As well as long-term evaluating what the damage is and what you need to do to -- for continuity of business, whether that be reputational or for for your -- your cyber program.

The next slide. Doing the little things right. You know, these are the things that, once again, that we recommend. And they toggle back and forth

between, as I've said, depending on what industry you're in, who your regulators are, whether -- or what country you're in or what state you're doing business in, whether or not these would apply to you is things that you would have to do or whether they're nice to do.

And as Kevin Greenfield said before, it's about harmonization, right? And as Doctor Turner Lee talked about soft law versus hard law. You know, that definitely, you know -- you know, rang a bell with me when she talked about that.

You know, sometimes I tell people, you really should do this. And sometimes I tell people you, kind of, have to do this. Because your regulators are going to come in and they're going to hammer you for not doing this.

Principle of least privilege. We talk about zero trust. This is an aspect of zero trust and as Dan Guido said before it's about minimizing the blast radius. You know, the threat actors are going to get in.

And if they get in and were you the one that

mentioned about Coca Cola? About hacking into Coca Cola. Getting into the EA and then getting the recipe for Coke. Not New Coke, just Coke. Right. Yeah.

Getting -- getting that recipe for Coke. You know, that's -- that's following the principle of least privilege, you only have access to what you strictly need to do your job. You do not have access to anything else.

This way when threat actors get in, they really can't move laterally as much. They can't get as much and it helps to limit the damage just like -- like Dan said before.

Multi factor authentication. You know, literally billions of breach passwords are out there. I see this all the time with, like, breach passwords being out there, being run back against your system through a password spray or credential stuffing attack.

Multi factor authentication, whether that be as we talk, it's -- it's something you know. Your password, obviously, something you have. Your PIN

card possibly. Or something that you are. Your thumbprint or biometrics.

And my conversation a couple hours ago with Dan was telling me this is all going out the window anyway within a couple of months, so, anyway. So my presentation is somewhat dated. But right now, this is something that you can do to tighten up your system.

Third-party vetting. You know, we saw before 53 percent of companies have gotten hit through somebody who came in through the third-party. And I can't tell you the number of times where I've had this happen where a client has said to me, I I think we got hacked.

And I'm like, well, you got victimized but you didn't actually get hacked. The hack came in through one of your third parties. And you can see this from smaller companies that are the third party may not be doing as much to lock down their cyber systems.

So obviously, much more due diligence needs to be done on the third parties that you're bringing

in because, as you said before, where you're bringing -- you're outsourcing the activity, you're not outsourcing the risk or the reputational damage. So that third party vetting in something that's actually key to -- to locking down your system.

And then lastly, MDR, Manage, Detection, Response. You heard a lot about a lack of skills out there. Not a lack of skills just a total demand for cyber skills outstripping who's actually out there. Smaller companies not having folks they need in house to help lock down their systems.

You know, that's not an excuse. Outsource it. You know there's talent out there and it -- and I don't do MDR, do I'm not, like, pitching my company here. Just outsource the MDR because then you're tapping into that pool where there's experts out there who are looking at what other companies are getting hit by and they can help you detect and block and protect your system.

So these are just an overview of some of the things that, as I said, I'm seeing here. And as I

said, toggling back and forth depending on the industry and between things that are good to do, nice to do and things that you absolutely have to do. But now I'll turn it over to Hilary for the regulatory piece.

MS. ALLEN: Thank you. So Tim's just discussed some important tools for addressing cyber security risk. I'm going to broaden the aperture a little bit and ask some bigger picture questions about how we think about and regulate operational risk.

So the points I'm going to make today are very much informed by the work that exists on understanding complex systems, which Dan already just alluded to in his presentation.

So complex systems are prone to cascade failures. So in cascade failures it's, sort of, the trigger is often unimportant. You can't predict the trigger but then once the trigger happens, they're unexpected interactions between the components of the complex systems with those interactions magnifying the problem as the failure cascades through the system.

So the whole system can fall down like a line of dominoes or in what's known as an overload failure, some components of the system can keep working but in doing so, they transmit problems to other components of the system, which can then fail.

So the robust, yet fragile dynamic that you sometimes hear about is critical to understanding a systems susceptibility to these kinds of cascade failures, which can be disabling. So a system can be robust in some ways but if too much emphasis is placed on the wrong kind of or just one kind of robustness and the others are sacrificed then the whole system as a whole is going to be more fragile to these kind of cascade failures.

So for example, I'm making this a little more concrete now. Efficiency is a kind of a robustness. Right? You want an efficient system. But if you put too much emphasis on efficiency, then the systems also going to be very good at moving around cascading problems as well.

And if you focus too much on efficiency then

you're probably also sacrificing the dimensions of robustness that help the system continue to function well in unanticipated circumstances.

So things like modularity, the ability to cope with changes in the organization of the systems components. Or scalability, the ability of the system to cope with changes to its own size and complexity.

Or evolvability, the ability to cope with changes to the systems usage over time. So if we think back to 2020 and the Covid pandemic, we learned this lesson with supply chains. Right? Steps that had been taken to make distribution more efficient in normal times left supply chains brittle when changes occurred.

Now, there is more and more interest in making components closer to home for supply chains that's less efficient but it's more robust. It recognizes, as Dan mentioned earlier, that sometimes this kind of systemic failure is inevitable. And so, you want some redundancy built in as well.

So when it comes to financial institutions

operations, I think it's time to start asking what some might consider a heretical question, which is, when is something efficient enough, such that making it more efficient will only damage the resilience of that system or the broader financial system.

So that's the first question I want to talk about is efficiency versus robustness trade off. And that has big implications for automation at large, the extent to which we automate things. We talked about this in the context of smart contracts but of course, that's not the only way to automate a system.

Next slide, please. So another provocative and big picture question I want to ask is, are cyber attacks really the biggest security threat that financial institutions technology systems face?

So I want to be clear here. I think cyber protections -- are tremendously important. But, as I said, the complex systems, sometimes the trigger can come in different forms and then cascade through the system.

And I think extreme weather events and even just glitches can have similar impacts to cyber attacks on how technology systems work. And they should be part of that same conversation as cyber often, yet they receive less attention.

Um, so by some estimates tech glitches, which is a non-technical term that I use to describe tech problems that happen accidentally rather than as a result of malicious actors. By some estimates, these tech glitches are more costly than cyber-attacks. Right?

The trigger is some kind of mistake, like a fat finger error or a coding error or something like that. Many financial institutions, I think, are particularly vulnerable to these kinds of glitches because their systems are so complex. They have coupled together multiple legacy IT systems and that makes them vulnerable to these cascade systems.

So just to give you a, sort of, an example of the types of things I'm worried about. So you can have a cascade failure if your support system or

automation systems respond to an initial problem and then accidentally introduce an additional failure as part of their response.

Or you could have a human trying to respond to a problem to mitigate or resolve a failure. But then their actions can lead to more failures. Often the response, and I think we all know this, to a compromised component is to reboot it, right? You turn it off and turn it back on. But the restart process can sometimes overload or otherwise throw off linked system components.

So we need to be thinking about these tech glitches that are, sort of self-inflicted harms in many ways, as well as the cyber attacks that come from nefarious actors. With regards to the impact of climate change on financial services, I think most of the discussion of climate risk that we hear in the financial regulation context focuses on the credit and market risks arising from the physical and transition risks associated with climate change.

But I think inadequate attention is being paid

to the potential for operational problems arising primarily from physical risks like extreme weather events that force office closures or knock out electrical grids or telecommunications lines that banks rely upon.

So all of these kinds of problems are becoming increasingly likely in this day and age. But their precise form and their impacts are uncertain. So we don't really have any historical data that will be predictive here.

And models that rely on the past are not going to do a good job of predicting operational risks in the future. So this, sort of, brings back to my earlier point of the need for redundancy in these spaces.

So the final slide, my final observation that I want to highlight here is when we think about systemic risks arising from operational problems in the financial system, we tend to focus on how financial losses from these operational problems might spill over.

So for example, a bank suffers a massive loss,

therefore it defaults on its loans to another bank, right? But that's a very real problem but it does miss the possibility that operational problems could be transmitted from institution to institution through technological channels.

And this picks up, I think, to some degree, on the third-party vender management. But I think it's a -- it could also, I think, increasingly as we see banks become linked to one another, either directly through APIs or through APIs with third parties who also interact with other banks, we have the potential for operational failure to invest in operational resilience.

And one bank, potentially, having systemic consequences for other banks, purely through technological channels rather than the typical financial spill overs that we're thinking about. However, we seem to be stuck in the mindset that operational risks are going to be idiosyncratic to the institutions that experience them.

So if you read something like the principles for financial market infrastructures, the PFMIs,

they focus on individual banks managing their operational risks within their own risk tolerances.

And I think you can draw analogies here with the micro-macro prudential discussion that we had after 2008. We used to think that as long as banks manage their own credit and market risks, whole system would be safe not realizing that the steps they might take to save themselves could impact someone else.

And I think we need to have that same macro-operational discussion or framework as we think about operational risk as well. Basically, it's possible, you know what I'm thinking the example that I think of here is, sort of, reliance on cloud infrastructure.

If everybody's trying to get to the same cloud and download their backups at the same time, the cloud might go down. Right? So I think we need to think about how steps banks may take to improve their own operational resilience could undermine the operational resilience of the system more broadly.

There's no reason to think that the cascade failures that I mentioned are going to respect a bank's organizational boundaries. Cascade failures could be a channel for bringing operational risks into the banks from outside and vice versa, as I said.

The key takeaway here is that decisions that financial institutions make about operational risk can impact other institutions as well. So the takeaway, as I said here, is that just as we -- sorry -- at the end of the slide.

Just as we realized after 2008 that we needed a macro prudential focus that took systemic interaction seriously with regards to credit, market and liquidity risks, we now need a macro-operational focus that extends this mindset to the transmission of operational risks.

So I know that's a lot of, sort of, really big picture questions but I think I'm thrilled about having the opportunity to work with this committee on thinking outside the box on these operational issues. So I just want to throw open these big

questions as a starting point.

MS. HOUSE: Thank you so much for the great overview of the threat landscape, Tim. And then Hilary, your discussion of macro-operational risks that don't respect boundaries. I think that has some really interesting implications for the Commissioner sitting next to me.

Then for our third and final presentation Steven Silberstein, CEO of the Financial Services Information Sharing and Analysis Center, the FS-ISAC, will present on the state of financial sector defense and collaboration to combat cyber threats. Steven over to you.

MR. SILBERSTEIN: Thank you. First a housekeeping question. Are we keeping to time? Because I can do an abbreviated version if that would be helpful.

MS. HOUSE: I -- honestly, we're interested in your presentation but if there is a briefer version -- I honestly defer to your judgement. We're looking forward to hearing about FS-ISAC.

MR. SILBERSTEIN: Well, given that it's at the

end of a -- of a jam packed session I'm going to forgo the slides, I think, and just stick on some content and try to do it in 10 to 12 minutes. Leave a few minutes.

MS. HOUSE: Perfect, thank you so much, Steven.

MR. SLIBERSTEIN: Okay. Thank you. First, thank you to Commissioners, Goldsmith Romero and Johnson and Chair House and Anthony, you have so many titles. Thank you for your work in the technology analysis committee and the CFTC.

I'll add that some of my perspective is from an operator where I've been most of my career in, principally, global capital markets, and I've had the honor of having a Series 3, 7 and 24 registration. So I've lived in the world.

I want to share first the industry perspective around the sectors security having -- in this job actually seeing a lot of sectors now. And I think we have a unique advantage in this rapidly evolving area.

And that is, there's both a fiduciary culture

of security as well as a regulatory culture stroke requirement. The fiduciary responsibility, it's bad to lose your customer's money. It goes back to having the strongest safe and the best guards with the biggest guns as a selling point to the bank. And it's continued now into a virtual and digital world as a necessity.

Also want to note, and I'll talk about it more that collaboration in the sector has a long culture of a public profit and also then the public sector within the private sector that goes back to the 1998, in the form of what was, became the Sector Risk Management Agency and ISAC.

And I think it exemplified here in that Ty Conklin, head of OSIP Treasury and the Sector Risk Management Agency, Kevin Greenfield, OCC and Dan, Trail of Bits, the organizations and the individuals all are engaged. They're not new faces to each other. And we're not talking about policy. We're talking about operational security on that.

I also want to throw in in my limited time, a little quip and a important -- a lot of what we're

talking about here in cyber defense and security is some degree base cyber hygiene, which is the most important thing.

And I think many folks will say regulation itself doesn't determine good cyber hygiene. It's the practice. An interesting model I use to test an organization's state of cyber hygiene is to find out if the CEO and the chair of the board have multi-factor authentication on their personal email. That gives you a good sense of the cyber awareness of an organization.

A little bit on FS-ISAC. We are in our 24th year. Founded in 1999. It is a member-driven, not-for-profit organization with a member-based board.

Our goal is to advance a cyber security and resiliency of the global financial sector. And principally we're protecting society's assets and the institutions that serve them. It's a cooperative in the truest sense.

It came from the early founder, who realized cyber security isn't a solo sport. And they were only as strong as the collective. And only as

strong as the weakest. So the organ -- sector got together to do that.

We really act as a force multiplier for the collective intelligence and somewise emphasize that you could be on a daily transactional basis and view us as a sensor network with over 5,000 members globally. Now many of the institutions are small and not necessarily contributing a lot but they also are importantly consuming and as is well, better able to defend themselves.

That trusted community's important because it knows how to respond both in protection and a response and some of the longer issues. To -- give you a sense of scale. Our members represent over $100 trillion of assets.

We currently have about 5200 firms in the -- as members. And 22,000 active users on our various intelligence and sharing platforms. We operate with three pillars of focus. Of course, cyber intelligence. And that's not just sharing of detail. We do a lot of enrichment. We do deep dives. Do topical sharing.

We -- overall security. How to maintain the security, best practices, white papers, webinars, et cetera; and resilience exercises are important part of what we do. Some of our own, some in collaboration of partners like FSSCC and FBIIC for the Hamilton series. And we'll discuss a little more about that.

You know, resilience, let me -- I'm going very topical for a moment. We've had two recent incidents. One was the anonymous Sudan threats and a real -- and that was threats and some bits of incident.

But one which has been pervasive through the financial services supply chain has been the MOVEit file transfer vulnerability. And every day we seem to learn about another firm in the supply chain who has lost some information.

This is a very interesting one because we haven't seen -- it hasn't been disruption. It hasn't been ransomware. It hasn't been any classic malware. It's been a huge exfiltration of PII from numerous sources. What happens with that? We don't

know yet.

What does it mean? Well, minimally it can be tremendous fodder for the types of attacks that Tim was just describing as far as identity theft, account logons, account compromises.

We'll see what happens. But it's also been an all-hands-on deck in much of the sector to understand who’s affected how. How to mitigate. Get some sense of the losses and how to potentially protect them further on. And we've been very actively engaged in that.

The -- one of the things that drive us in the trusted community is that the traffic light protocol, which we created many, many years ago. It's basic information handling rules, which is broadly used in this sector and variations in many other sectors that allow you to have a trusted mechanism to say what am I sharing, with who?

Example, if it's TLP-Red and I'm having discussion with Carole, that means it goes no further than her. If it's TLP-Amber it can go to the people who need to know in the organization.

And then it goes from there. That's the key part of it.

And I'll add a few more things in the intro. We're not unilateral just in cyber. We have two very interesting subsidiaries. We have Sheltered Harbor, which is focused on, in simplest terms, a protocol and standard that any firm can implement to protect their own operational data from a ransomware or operationally destructive incident by creating an archival protected, encrypted, air gapped and recoverable copy of key -- well, any data -- but key operational data.

We also have FDX, the financial data exchange, which has created the standard for open banking data interchange, which is currently supporting over 60 million consumer accounts between financial institutions and the open banking world.

Very relative to this community -- and a little aside here -- 5200 firm members is a big number. We don't treat it as one big world. We have numerous communities of interest with leaders in FS-ISAC to communicate, promote, organize that

match sector, subsectors and geographies.

So we have a long standing in the securities industry risk group. Try to give them some -- some good names, which is composed of 450 firms including FCMs, broker dealers, asset managers, retirement firms, and alternative investors with subgroups within that.

Over 1000 individuals participate in that community in various forms part of it from real time intelligence to more strategic work. Adjacent to that, also very relevant, is the clearing house and exchange form, the CHEF for short, which has 22 clearing houses and exchanges globally, which represent -- which virtually all the major exchanges and clearing houses around the world.

This is very strategic. We have been very engaged with them and two relatively recent examples. Of course, the most recent one being Ion Trading and that outage. And previously, a big public event was New Zealand Stock Exchange outage due to a DDoS. And that was about, at this point, three and a half years ago.

So we are in the middle of the space. Approximately, 60 percent of the US FCMs are members of FS-ISAC. Let me talk in these few minutes about what we see as the threat landscape going forward. And I think the previous presentations did a great job in describing it. So I'm going to be a little more generalized and go from there.

The challenge tomorrow is that today's vulnerabilities are different than tomorrow's. Today's techniques, TTPs are different than tomorrow's. And what is leading edge today has to become tomorrow's baseline for a firm to protect against and be able to respond against. And that is an ever-increasing burden.

Our attack surface is ever expanding because as Kevin noted the proliferation of important service providers increases as we become more virtual outside the firewall of the institution in that.

And I leave it to the Commission to come up with interesting ways to deal with that. We are

doing a few things towards that, which include working with the FSSCC/FBIIC cloud initiative or in one of the work streams around transparency and resiliency.

We have created a critical provider program for some of the large missing critical infrastructure operators with whom the sector is dependent on to get work done every day. That's cloud service providers and related type firms.

That's a little aside. Let me come back to the other -- more challenges. We are seeing that the -- and MOVEit is a great example of the following. The required response time for published vulnerabilities is getting close to zero. With the technology to scan the public internet and find the vulnerabilities is great to protect but it's also great for the adversaries to jump on it.

So having better response, not waiting to patch but turning off vulnerabilities is going to be key. But also, result in business interruption as a result. We're not going to see, as noted, ransomware. Ransomware is a service. Malware is a

service. It's not going away. The ROI is still positive, and we continue to see the tactics change.

Even DDoS, which was old -- ten-year-old story in this most recent anonymous Sudan work, we saw some interestingly evolve -- interesting evolving tactics that we weren't as protected against as we should.

All that said, the number of serious intrusions into financial sector firms, the regulated part of the world, are still fairly low on any measure. But that doesn't mean we should be satisfied because I know the protections need to continue.

As we look more strategically, and I originally wrote one to five years out but some of this feels even more real. In a world of quantum computing, it will never be too early to start addressing the gargantuan task of coming up with post-quantum capable photography because we have to rewire a very complex sector not unilaterally, like we did with the Y2K, but we have to do it well or

our transactions are airplanes that are in flight.

We recently published five papers from our PQC working group, which are on our website, which I -- is really about assessing, understanding and starting the plan. The idea being know where your assets, cryptographic assets are today because they are all over the organization and be ready for it.

The -- a lot has been heard about AI and generative AI. I -- when I said one to five years, the longer-term concern is that identity as we know it, our image, our moving image, our voice, even potentially some of our behavior becomes at risk. I joke that we may require people to come back into the branch to prove who they really are because that's going to be a challenge.

And we're already seeing generative AI break down some of the language and geographic barriers that have been protecting places. So western languages tend to be easy to simulate for phishing and smishing. It was hard for much of the world to deal with Arabic languages and eastern Asian character-based languages. That obstacle is now

gone.

Anybody can essentially do a reasonable translation and suddenly I can be phishing and smishing in Japan. That is real, it's being felt; and it's not going away quickly. And there's no regulation that is going to be able to protect that because these tools will always be available to the bad guys.

We also have concern, long term, as we are doing a good job to increase end user security, end user awareness out in the consumer base. What does that do for the digitally challenged? It's become even more of an obstacle for them.

One of our biggest areas of concern is around the things like retirement accounts and senior citizens who are already somewhat digitally challenged. We need to better protect them but they're not adept with the toolkit. So concerned about that.

Much has been mentioned about LLM model pollution. I want to add to that. In general, I'll give a simple real world. We already suffer from a

little too much of, I saw it on the internet. There it must be true. So it must be true.

We face the same issue with I saw the ChatGPT said, therefore it must be true. Relative to keeping confidence in our institutions, confidence in a system, which as a whole works fairly well, we have a longer-term mis, dis, mal information MDM challenge around the AI tools.

The -- and I mentioned before, I continue to mention this is the challenge for the smaller institutions to have the budgets, the manpower, the talent to secure the organization where there's an ever-increasing baseline requirement for that core cyber hygiene need to be -- that needs to be present in financial services.

Quickly just note some of the areas where the sector has collaborated broadly across cybersecurity and resiliency; numerous FSSCC collaborations. Thank you to everybody involved.

We have a sector-wide all-hazards playbook supported by something called the core executive response group, which -- not known by many, because

we haven't really published it, was publicized the fact that January 30th, 2020, this sector turned up the playbook in their response mechanism for Covid-19.

It may have been the first private -- well, the first in the west, one of the first efforts turned up anywhere and it was amazing cooperation that resulted in the finance sector basically continuing to operate fairly well in critical transactions going on through Covid-19.

Wasn't a traditional cyber event. Wasn't an incident when we called for this to happen. It was a, hey, some incident may occur, and we need to be prepared. And a fascinating effort. That effort continued almost non-stop because of many other incidents and including SolarWinds and Russia-Ukraine. It's now quiet because more stable world but we have the ability to spin that up.

The Hamilton exercise, which is our FSSCC/FBIIC collaboration, continues to look at the hard problems, not the easy ones. It's not an exercise about muscle memory. This is an exercise

about what if this really important thing broke? What do we do?

That was one of the generators of the FSI -- sorry, the Sheltered Harbor effort. What happens if we had a Sony Pictures Entertainment-style attack against a financial institution and all their data was wiped out including the backups? What's our last line of defense?

I noted the Ion Trading effort has been a very important point of collaboration to determine if it was a systemic risk or not. CFTC was a great contributor to that assessment and both in the public sector and the private sector that was both calming and also gave us a clear response. But we had a unprecedented calming message from FBIIC/Treasury that was extremely helpful in that situation.

We emphasize at [inaudible] a series of exercises. We're shortly starting our annual CAPS exercise. We kept the name because it works -- Cyberattack Against Payment Systems. But we now do variants of it that are sector specific --

subsector specific, insurance, banking and securities.

Start off with a common attack point but then it goes into a business problem specific to the sector. Interesting technology. It reaches -- last year, 10,000 people, 1,000 firms. So it is a distributive tabletop, semi-synchronous -- very powerful mechanism.

We have numerous working groups, which are trying to work both on the tactical and the leading edge. Some of the leading-edge points are an AI working group. Not trying to reinvent the wheel but to look at a framework for firms to evaluate their usage. Not the answer but the mechanism, as well as I mentioned the PQC piece.

In the spirit of time, supply chain. I hit in numerous factors, I'm going to leave it there and just add that we have a business resiliency committee, which is focused on the what's after. And again, this is member-driven because if you have a serious cyber incident, it's not a technology problem; it's a business problem.

And a pitch that we're trying to do our little piece in the talent development challenges in the finance-cybersecurity sector with a scholarship program. So probably went a little extra. Sorry for that.

MS. HOUSE: It was a great presentation. Thank you so much, Steven. Great as always to hear from the FS-ISAC. And to hear about the threats presented to sounds like the majority of the FCMs under the authority of the Commission and as well as the collaboration efforts underway.

We're going to have a very, very brief discussion because I know Nikos, you had raised your flag. And then Jonah, if it's okay, I'd love to turn to you, which may be the final question unless someone else is flagged or the Commissioner has a comment. But I think that you spoke recently on third-party risk items on a podcast. So if it's okay, after he goes, I'd love to turn to you for any remarks from yourself.

MR. ANDRIKOGIANNOPOULOS: Yeah. I just want to make a very quick comment on third-party risk

management and connect that with decentralization.

I think one of the observations that I have is when it comes to financial institutions dealing with decentralized services, decentralized is the last services, I think some of the risk that we're seeing is that some of the larger financial institutions understand the compliance and the monitoring that they need to be doing.

And they do that for reputational reasons even though the regulators are not asking for it. But even in those cases, but smaller ones, smaller financial institutions are much more willing to enter into third-party agreements where they don't actually realize where their responsibility and the risk, kind of, lies.

I think that calls for decentralization kind of rules and regulation become clear sooner so that everybody realizes and doesn't rely on those third-party relationships and the contracts cover those risks.

MS. HOUSE: Kevin, is that a reaction to Nicol's comment?

MR. GREENFIELD: Yeah. No. That actually is an excellent point. And one of the things we had done previously is actually putting out for community banks a due diligence guide for engaging with FINTEC organizations because, you're absolutely right.

You have some very large technology service providers that are used to operating within the financial sector or used to operating with government agencies that understand expectations and controls.

But a lot of these emerging companies don't have that experience and knowledge. And then when you have less sophisticated firms that engage with them, it can run into problems. So one of the things we did was highlight what are some unique ways to approach due diligence activities with firms that don't have a 50-year history of operation.

That don't necessarily have all the same types of reports that you can request that some are more sophisticated. So it is a key issue, and it is

something because we definitely have seen where it has cost issues in the industry.

MR. ANDRIKOGIANNOPOULOS: And I think, Kevin, your point is that the FINTECs understand decentralization, but the financial institutions understand compliance. So as they partner there needs to be some common understanding between the two.

MS. HOUSE: All right. Jonah?

MR. CRANE: Sure. So for my sins, thank you. Thanks for the great presentations. I mean, I was struck. I'll just be real quick in making a couple observations.

I was struck in the conversation today that almost everything we talked about is going to be looked at by financial institutions through the third-party risk management lens because they're relying on third parties to provide all kinds of technology solutions.

And increasingly, it provides the solutions that are helping them oversee the technology solutions. Right? So I think Commissioner Pham

mentioned model risk management in the beginning.

We talked about the need to monitor for fraud manipulation in using high frequency trading algorithms. We talked about the risk of bias in using some of these new algorithms. We talked about how to oversee AI generally and whether overseeing the models or, sort of, the governance was the right approach.

There are start ups out there and really established firms out there trying to provide solutions for -- to oversee every single one of those risks. And so, now you're into third-party and fourth-party and fifth-party land, as Kevin said.

I think all of this really -- Hilary's presentation, sort of, summarized my thinking probably most closely, which is the supply chain around financial services and financial markets is just increasingly complex.

This is not necessarily new, but I think it's -- I think it's becoming really complex in new ways and really raises the risk of the kind of cascade

failures that Hilary talked about. And ultimately, the burden in our current system falls on the regulated institution.

I think Dan's presentation really highlighted nicely how that sort of system of delegated responsibility raises real questions and conundrum when you think about applying it in the decentralized world.

Kevin emphasized the importance of contractual requirements. And Nicole, your question just now highlighted to me that if that's one of your primary protections, how do you even think about that in a decentralized world.

I actually think one of the interesting things that we should spend more time thinking about is what role might be played for, sort of, more collaboration and cooperation and maybe use of utilities might be able to be leveraged.

Both for the regulated, sort of, current trapped by sector where, I think, the complexity of all this might become overwhelming but also in a decentralized space where maybe you have that, sort

of, third-party validation that folks can be -- start to rely on.

And you don't have all the protections that are laid out in the recent guidance but maybe -- maybe it's something that we can leverage going forward. I know there have been efforts in the past and they've not become widespread for various reasons but the need for them may be -- may be greater now. So I'll stop there.

MS. HOUSE: Great observations. Thank you so much, Jonah. I understand that Steve has a -- has a question from our Zoom participation.

MR. SUPPAN: Yes. This is for Mr. Silberstein. Thank you for a very sobering and yet somehow comforting presentation. I've never heard about the all hazards playbook. Have you applied the all hazards playbook to the eventuality of a green swan event?

There was a study by the United Kingdom's Association of Actuary Accountants that said a current global bank climate financial risk preparation or modeling is implausible because the

actuarials are so insufficient relative to the likely costs and sequence of time related tipping points. So that's my -- my question.

MR. SILBERSTEIN: I'm going to describe the playbook as a framework for a response and collaboration and acceleration. It's not a detailed what if the following happens because the sector is way too complex both in what's known and, as you're mentioning, what's unknown to have a descriptive playbooks for everything.

There are some very distinct ones around payments, which was worked on with the ARC and a playbook which is managed by SIMA. There are various FIBIC distinct playbooks but there is nothing -- we can't be overreaching.

Most importantly, it's a framework to say who calls who, when, and how do we engage in being adaptive in the -- to some degree, it's an adaptive battle plan versus a prescriptive answer. I hope that answers the question.

MS. HOUSE: Thank you -- Oh, sorry, yes, Todd.

MR. SMITH: I'm Sorry. Yeah. Just really

quickly I mean, and I -- Steven had mentioned Ion and MOVEit as the two most recent examples where we took those playbooks and actually implemented them.

And it's a real time exchange so for MOVEit, I was alerted to MOVEit around 9:00 p.m. on a Friday evening. We, at 9:01 we activated the treasury instant response playbook, which means the secretary and deputy secretary immediately get notified that there's a -- there's an incident that maybe have some -- have some touch points in the financial sector.

Two minutes after that Steve Silberstein had headed that FS-ISAC and then also the head of the FISIC, Ron Green, who's the chief security officer. A measure gets notified that we've activated the playbook. Kevin Green in the field gets notified five minutes later.

So you're talking about afterhours on a Friday evening while we still haven't even developed the common operating picture, everyone starts to come together. And then we coalesce around developing a common operating picture.

By Saturday morning we already had a fair amount of lead already on what we thought the worse case scenario was going to be. And we worked that throughout the weekend and in lock step with one another. It's a true public private partnership and the communication flows in both directions.

MS. HOUSE: Thank you so much Todd. Great to hear from the SCMA about the value of collaboration and timeliness of doing so for the financial sector. It is now time for closing remarks from Commissioner Goldsmith Romero.

MS. GOLDSMITH-ROMERO: It's so interesting that you said that Todd because I remember that Friday night because I was constantly getting updates from our SESO on the -- on MOVEit and what was happening in our markets which was fascinating to, kind of, watch. And it continues to spill out and we still don't know, sort of, the event of that.

Very much appreciate all the perspectives and Hilary, on your operational resilience and operational risk and thinking about it broader, I

think that's, sort of, the way to do it.

We are moving, I think, from thinking about just cyber to operational risk. We are also moving from a mentality of incident response to one of resilience and building resilience, that idea of bouncing back from setbacks.

And this gets to what Tim was talking about, which is the idea of plan for it and then test it and then try it out and work it out and monitor it. And then the idea from a government perspective, thank you Kevin for coming in.

This idea of trying to build in something that is flexible and can also stand the test of time but is also proportionate and commensurate with the risk involved, which -- which hopefully then I think is the point you were getting at. Sort of scales to basically the size of the institution or company, as well as the type of risk, the type of provider.

I think these are all the types of things that we have to consider. And just before this meeting we were also talking, in my office, with some of

the -- with the chairs and the co-chairs about AI enabled and also Quantum and the move from that.

So thank -- thank you for that. And I very much appreciate everyone involved. All of the presentations today were wonderful. We could discuss these for hours. The most important thing I am -- I'm really liking is the idea that this group is coming together to have a trusted environment where you feel comfortable raising your views and adding your viewpoint.

And as the slides are coming in and as the discussions are happening, I'm watching that. The diversity of viewpoints is very, very beneficial to the CFTC. And I think very beneficial to a number of policy makers out there.

I've heard a lot of others that are watching our work, that are watching these -- others will watch this later. And I get a lot of comments on it and I'm very, very grateful, ultimately, for your public service. This all goes into the idea that, in my mind, government must keep pace with technology, or the most vulnerable people will

suffer.

How we do that is to bring all of you in with your different expertise, your different viewpoints and share -- I appreciate you sharing your experience, sharing your broad viewpoints. And for that I'm truly grateful.

Sorry for going over in time but you know we could just keep going on this conversation. Thank you again, Kevin. Thank you, Steven. Thank you for all of our presenters today and all of our TAC members.

And there's important work to be done in the subcommittee so plan on rolling up your sleeves. Very much appreciate this is not your day job but I view this as public service. I appreciate you answering that call to public service so, thank you.

And, of course, very big thank you to Chair -- our Chair, Carole House and Tony and Lauren and Drew and Scott Lee in my office who worked very hard to put this agenda together thinking very deeply about the types of things you might want to

discuss.

Obviously, we have future meetings. So I'm sure you guys won't be shy. Thank you very much. I very much appreciate your public service.

MR. REDBORD: Meeting adjourned.