

A central graphic of a square chip with the letters 'AI' in white. The chip is surrounded by a complex network of blue lines and circles, resembling a circuit board or neural network. The background is a gradient of blue with various geometric patterns and lines.

AI

RESPONSIBLE ARTIFICIAL INTELLIGENCE

IN FINANCIAL MARKETS:
OPPORTUNITIES, RISKS & RECOMMENDATIONS

*A Report of the Subcommittee on Emerging and Evolving Technologies,
Technology Advisory Committee of the
U.S. Commodity Futures Trading Commission*

A central graphic of a square chip with the letters 'AI' in white. The chip is surrounded by a complex network of blue lines and circles, resembling a circuit board or neural network. The background is a gradient of blue with various geometric shapes and patterns.

AI

This report has been approved by the Subcommittee on Emerging and Evolving Technologies (Subcommittee) of the Technology Advisory Committee's (TAC). The views, analyses, and conclusions expressed herein reflect the work of the Subcommittee, and do not necessarily reflect the views of the TAC, the Commodity Futures Trading Commission or its staff, or the U.S. Government.

Reference to any products, services, websites, organizations, or enterprises, or the use of any organization, trade, firm, or corporation name is for informational purposes only and does not constitute endorsement, recommendation, or favoring by the U.S. Government.

Commissioner Christy Goldsmith Romero, Sponsor

Scott W. Lee, Senior Counsel and Policy Advisor, Office of Commissioner Goldsmith Romero, CFTC

Dr. Nicol Turner Lee, Co-Chair, Subcommittee on Emerging and Evolving Technologies, TAC

Todd Smith, Co-Chair, Subcommittee on Emerging and Evolving Technologies, TAC

Anthony Biagioli, Designated Federal Officer, CFTC

MAY 2, 2024

Table of Contents

Members of the Subcommittee on Emerging and Evolving Technologies on the CFTC Technology Advisory Committee (TAC)	4
Glossary of Terms	5
Executive Summary	9
I. Introduction	15
II. Defining AI and Responsible AI	20
III. The AI Policy Landscapes – U.S. and Abroad	26
Domestic Policies	26
Global AI Policies.....	33
Policy References to Financial Services Protections	35
IV. Financial Services AI Use Cases	36
Other Considerations.....	38
AI Applications Among CFTC-Registered Entities	45
Opportunities of AI Use Among CFTC-Registered Entities.....	46
Risks and Challenges	47
V. AI Risks and Management Among CFTC-Registered Entities	48
AI Risks	48
AI Risk Management Framework for CFTC-Registered Entities.....	52
Existing Risk Management Requirements	53
Paths Forward for AI Risk Management in CFTC-Regulated Markets.....	55
VI. Recommendations	58
APPENDIX ONE: ANNOTATED RISK TABLE	64

Members of the Subcommittee on Emerging and Evolving Technologies on the CFTC Technology Advisory Committee (TAC)

Name	Title	Affiliation
Dr. Nicol Turner Lee (Co-Chair)	Senior Fellow and Director, Center for Technology Innovation	The Brookings Institution
Todd Smith (Co-Chair)	Director, Information Systems	National Futures Association
Dan Awrey	Professor of Law	Cornell Law School
Cantrell Dumas	Director, Derivatives Policy	Better Markets
Dan Guido	CEO and Founder	Trail of Bits
Carole House	Executive in Residence, and Senior Fellow, the Atlantic Council	Terranet Ventures Inc.
Ben Milne	Founder and CEO	Brale
Dr. Francesca Rossi	AI Ethics Global Leader	IBM
Joe Saluzzi	Partner, Co-Founder, and Co- Head of Equity Trading	Themis Trading
Dr. Steve Suppan	Senior Policy Analyst	Institute for Agriculture and Trade Policy
Corey Then	Vice President and Deputy General Counsel, Global Policy	Circle
Dr. Michael Wellman	Professor, Computer Science & Engineering	University of Michigan
Todd Conklin	Chief AI Officer, and Deputy Assistant Secretary of Cyber	U.S. Treasury Department

The Subcommittee also acknowledges the research support of Jack Malamud and Joshua Turner from the Brookings Institution, as well as Michael Schorsch from NFA.

Glossary of Terms

Artificial Intelligence (AI): “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.”¹

AI system: “any data system, software, hardware, application, tool, or utility that operates in whole or in part using AI.”²

AI safety: “an area of machine learning research that aims to identify causes of unintended behavior in machine learning systems and develop tools to ensure these systems work safely and reliably.”³

Algorithmic bias: “the systemic and repeatable errors in a computer system that create unfair outcomes, such as privileging one arbitrary group of users over others.”⁴

¹ National Artificial Intelligence Initiative Act of 2020, Pub. L. 116-283, div. E, § 5002, 134 Stat. 4523 (2021), 15 U.S.C. § 9401(3), <https://www.govinfo.gov/content/pkg/USCODE-2022-title15/pdf/USCODE-2022-title15-chap119-sec9401.pdf>.

² Joseph R. Biden Jr., “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” October 30, 2023, p. 6, <https://docs.house.gov/meetings/FD/FD00/20240206/116793/HHRG-118-FD00-20240206-SD001.pdf>.

³ Tim G. J. Rudner and Helen Toner, “Key Concepts in AI Safety: An Overview,” Center for Security and Emerging Technology, March 2021, <https://cset.georgetown.edu/publication/key-concepts-in-ai-safety-an-overview/>.

⁴ Abid Ali Awan, “What is Algorithmic Bias?,” DataCamp, July 2023, <https://www.datacamp.com/blog/what-is-algorithmic-bias>.

Deep learning: a subfield of machine learning that uses neural networks with many layers (“deep neural networks”) to learn complex functions.⁵

Deepfake: “a video, photo, or audio recording that seems real but has been manipulated with AI.”⁶

Foundation model: “any AI model that is trained on broad data... that can be adapted to a wide range of downstream tasks.”⁷

Frontier model: a “highly capable general-purpose AI [model] that can perform a wide variety of tasks and match or exceed the capabilities present in today’s most advanced models.”⁸

Generative AI: “the class of AI models that emulate the structure and characteristics of training data to generate derived synthetic content. This can include images, videos, audio, text, and other digital content.”⁹

⁵ Samuel K. Moore, David Schneider, and Eliza Strickland, “How Deep Learning Works,” IEEE Spectrum, September 28, 2021, <https://spectrum.ieee.org/what-is-deep-learning>.

⁶ Government Accountability Office, “Science and Tech Spotlight: Deepfakes,” GAO-20-379SP Deepfakes, February 2020, <https://www.gao.gov/assets/gao-20-379sp.pdf>.

⁷ Rishi Bommasani et al, “On the Opportunities and Risks of Foundation Models,” Stanford University, July 12, 2022, p. 3, <https://arxiv.org/pdf/2108.07258.pdf>.

⁸ UK Government Office for Science, “Future Risks of Frontier AI,” October 2023, <https://assets.publishing.service.gov.uk/media/653bc393d10f3500139a6ac5/future-risks-of-frontier-ai-annex-a.pdf>.

⁹ Executive Order 14110, p. 8, *supra* n. 2.

Hallucination (AI): the phenomenon in which an otherwise capable large language model presents false information as fact or generates nonsensical output that is unrelated to the prompt.¹⁰

Machine learning: the dominant subfield of modern AI in which algorithms use training data to teach themselves how to correctly make predictions, analyze data, or perform tasks, as opposed to methods in which humans explicitly program the behavior of the system.¹¹

Red team testing: the practice of stress-testing a model to find failure modes, such as prompts that bypass content guardrails in a large language model.¹²

Reinforcement learning: a subfield of machine learning that aims to design agents that act optimally in a real or virtual environment to achieve some goal by maximizing their expected “reward”.¹³

Smart contract: “a term used to describe computer code that automatically executes all or parts of an agreement and is stored on a blockchain-based platform.”¹⁴

¹⁰ Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung, “Survey of Hallucination in Natural Language Generation,” ACM Computing Surveys 55, no. 12 (March 3, 2023), 1–38, <https://doi.org/10.1145/3571730>.

¹¹ Sara Brown, “Machine learning, explained,” MIT Sloan Management School, April 21, 2021, <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>.

¹² Jessica Ji, “What Does AI Red-Teaming Actually Mean?,” Center for Security and Emerging Technology, October 24, 2023, <https://cset.georgetown.edu/article/what-does-ai-red-teaming-actually-mean/>.

¹³ “Reinforcement learning,” GeeksForGeeks, <https://www.geeksforgeeks.org/what-is-reinforcement-learning/>.

¹⁴ Stuart D. Levi and Alex B. Lipton, “An Introduction to Smart Contracts and Their Potential and Inherent Limitations,” Harvard Law School Forum on Corporate Governance, May 26, 2018, <https://corpgov.law.harvard.edu/2018/05/26/an-introduction-to-smart-contracts-and-their-potential-and-inherent-limitations/>.

Supervised learning: a major paradigm of machine learning that aims to make predictions based on labeled training data, in contrast to unsupervised learning.¹⁵

Training data: the dataset from which a machine learning algorithm learns statistical relationships to make predictions or otherwise produce outputs.¹⁶

Unsupervised learning: a major paradigm of machine learning that aims to find patterns in unlabeled data, in contrast to supervised learning.¹⁷

Untapped unstructured textual data: the wealth of text data in the world that has yet to be integrated into datasets and which could help further increase the capabilities of large language models.¹⁸

Value Alignment: the degree to which an AI system is “aligned not only with value-neutral human preferences (such as intentions for AI systems to carry out tasks) but also with moral and ethical considerations.”¹⁹

¹⁵ “Difference between supervised and unsupervised learning,” GeeksForGeeks, <https://www.geeksforgeeks.org/difference-between-supervised-and-unsupervised-learning/>.

¹⁶ Amal Joby, “What Is Training Data? How It’s Used in Machine Learning,” G2, July 30, 2021, <https://learn.g2.com/training-data>.

¹⁷ “Difference between supervised and unsupervised learning,” *supra* n. 15.

¹⁸ Bikram Dahal, “Unleashing The Power Of Unstructured Data: The Rise Of Large AI Models,” Forbes, July 24, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/07/24/unleashing-the-power-of-unstructured-data-the-rise-of-large-ai-models/>.

¹⁹ Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Kwan Yee Ng, Juntao Dai, Xuehai Pan, Aidan O’Gara, Yingshan Lei, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, Wen Gao, “AI Alignment: A Comprehensive Survey,” February 26, 2024, <https://arxiv.org/pdf/2310.19852.pdf>.

Executive Summary

The proliferation of artificial intelligence (AI) is impacting nearly every industry, and the financial services sector is no exception. In 2023, 99% of financial services leaders reported that their firms were deploying AI in some capacity, and these emerging and evolving technologies have the potential to support a variety of use cases related to financial services. In 2022, Commodity Futures Trading Commission (CFTC) Commissioner Christy Goldsmith Romero began sponsoring the reestablished Technology Advisory Committee (TAC) to focus on the intersection between technology, law, and policy. The Emerging and Evolving Technologies Subcommittee (Subcommittee), from which this report is delivered, convened to research and compile findings around the opportunities and challenges of AI adoption and use by entities under the jurisdiction of the CFTC, including but not limited to exchanges, clearinghouses, banks and others acting as futures commission merchants and swap dealers, as well as managed funds and advisors, introducing brokers, retail foreign exchange dealers, and data repositories. This group will collectively be referred to in the report as “CFTC-registered entities.”

This Subcommittee defines critical terms, such as “responsible AI,” “generative AI,” and “AI governance,” before proceeding with a comprehensive review of the current AI policy landscape. That is, we explore the landscape of research and public policies from the United States and abroad to contextualize the evolving nature of regulatory guidance and voluntary best practices using AI and offer other considerations for the CFTC as it explores how to address and mitigate local and global risks, including the AI Risk Management Framework presented by the National Institute of Standards and Technology (NIST). The report also offers a series of AI use cases as they currently or could potentially apply to CFTC-registered entities, including fraud detection and prevention, customer relationship

management, predictive analytics, and credit risk management, among other examples, and provides guidance on risk management in these and other financial markets scenarios.

In theory, AI represents a potentially valuable tool for CFTC internal and external stakeholders. The value proposition primarily resides in the potential of AI and other evolving technologies to improve automated processes governing core functions, including risk management, surveillance, fraud detection, and the identification, execution, and back-testing of trading strategies. AI can also be used to collect and analyze information about current or prospective customers and counterparties. However, despite its potential value, the use of AI by CFTC-registered entities will require further exploration and discussion, particularly raising awareness as to the function of automated decision-making models, and the necessary governance. Part of the challenge in crafting this report was the lack of direct knowledge about the CFTC-registered entities currently leveraging AI, and the level of transparency and explainability among these firms for regulators and customers about AI's use, particularly as it pertains to trading strategies and risk management. But even where firms disclose their use of such technologies, it is not always clear the type of AI that they are using (e.g., predictive, algorithmic, generative, or other frontier models), and for what use cases. Other considerations in AI use include responsible development, the quality of training data, the extent of involvement of humans in autonomous trading models, data privacy, auditing and oversight, and the breadth of internal talent at the CFTC or at registered entities to perform all or some of these suggested activities.

To begin the conversation on the design, use, and governance of AI models, this report from the Subcommittee begins with more granular focus on the opportunities and risks of increased AI applications and integrations into financial markets and the possible use cases among CFTC-registered entities. The report concludes with a series of proposed recommendations to the CFTC that explore the adoption of AI by their registered entities, and offer additional proposals around the Commission's

governance structure, including the use of its rulemaking authority, staff consultations, and other public means to make the adoption and use of emerging and evolving technologies much more transparent, and shared among regulated companies, as well as third parties. The Subcommittee finds that, without appropriate industry engagement and relevant guardrails (some of which have been outlined in existing national policies), potential vulnerabilities from using AI applications and tools within and outside of the CFTC could erode public trust in financial markets, services, and products. These factors can also make the supply chain less predictable. While there is still much more to be discovered and explored with the proliferation of AI among entities regulated by the CFTC, the Subcommittee suggests that the Commission and the specific role of the TAC start to develop a framework that fosters safe, trustworthy, and responsible AI systems, and consider the five *Proposed Recommendations* below.

Proposed Recommendations

A. Recommendation One: The CFTC should host a public roundtable discussion and CFTC staff should directly engage in outreach with CFTC-registered entities to seek guidance and gain additional insights into the business functions and types of AI technologies most prevalent within the sector.

- *Intended purpose*: To inform the CFTC about key technical and policy considerations for AI in financial markets, develop common understanding and frameworks, build upon the information of this report, and establish relationships for future discussion. The discussion topics should include, but not be limited to, humans-in-or-around-the-loop of the technology, acceptable training data use cases, and development of best practices and standards as it relates to the role of AI.

- Intended audiences: CFTC regulators, staff, registered entities, and third parties who can use insight to design appropriate AI regulatory policy and/or best practices.
- Potential outcomes/deliverables: Roundtable discussion and supervisory discussions and consultations with CFTC-registered entities to ascertain how AI systems are used in markets and how future AI developments may impact markets.

B. Recommendation Two: The CFTC should consider the definition and adoption of an AI Risk Management Framework (RMF) for the sector, in accordance with the guidelines and governance aspects of the NIST, to assess the efficiency of AI models and potential consumer harms as they apply to CFTC-registered entities, including but not limited to governance issues.

- Intended purpose: To ensure some certainty, understanding and integration of some of the norms and standards being developed by NIST, and to introduce these practices to regulated industries and firms.
- Intended audiences: CFTC regulators, staff, registered entities, and others who will be impacted by the AI RMF, and its intended outcomes.
- Potential outcomes/deliverables: A potential proposed rule from the CFTC applying the information from Recommendation One to implement the NIST framework, thus ensuring financial markets and a regulatory system that is more resilient to emerging AI technologies and associated risks.

C. Recommendation Three: The CFTC should create an inventory of existing regulations related to AI in the sector and use it to develop a gap analysis of the potential risks associated with AI systems to determine compliance relative to further opportunities for dialogue, and potential clarifying staff guidance or potential rulemaking.

- Intended purpose: To confirm CFTC's oversight and jurisdiction over increasingly autonomous models, and to make more explicit compliance levers.
- Intended audiences: CFTC regulators, staff, registered entities, and others who have interest in AI's compliance.
- Potential outcomes/deliverables: Issue clarifying staff guidance or proposed potential rulemaking from the Commission that advances the explicit and implicit applications of existing regulatory measures.

D. Recommendation Four: The CFTC should strive to gather and establish a process to gain alignment of their AI policies and practices with other federal agencies, including the SEC, Treasury, and other agencies interested in the financial stability of markets.

- Intended purpose: To leverage and utilize best practices across agencies, and potentially drive more interagency cooperation and enforcement.
- Intended audiences: CFTC regulators, staff, and other similarly aligned regulatory agencies.
- Potential outcomes/deliverables: Interagency meetings and cooperation.

E. Recommendation Five: The CFTC should work toward engaging staff as both ‘observers’ and potential participants in ongoing domestic and international dialogues around AI, and where possible, establish budget supplements to build the internal capacity of agency professionals around necessary technical expertise to support the agency’s endeavors in emerging and evolving technologies.

- *Intended purpose:* To build the pipeline for AI experts at the agency, and to ensure necessary resources for staff, events, and other related activities that ensure more responsible engagement of AI by internal and external stakeholders.
- *Intended audiences:* CFTC regulators and staff.
- *Potential outcomes/deliverables:* Increased budgetary resources toward AI services, and more presence at relevant conferences and convenings.

These recommendations are also presented with the intended audiences for each, and potential deliverables. In some of the recommendations, we urge the CFTC to leverage its role as a market regulator to support the current efforts on AI coming from the White House and Congress.

I. Introduction

The design and deployment of artificial intelligence (AI), including the more recent generative AI models, are occurring within many sectors, including the financial services industry. The use of AI models including statistical and machine-learning techniques has been long-standing and consistent among both institutional and consumer-oriented financial services companies. Before the proliferation of generative AI, which relies mainly on Large Language Models (LLMs), the “unprecedented availability of affordable computer power” and the enhanced accessibility of hardware used for high performance computing have enabled the development and deployment of more autonomous models. Emerging and evolving AI technologies have also encouraged firms to explore integrating such systems into their trading strategies and business models.²⁰

In 2017, a survey of executives conducted by Deloitte and the European Financial Management Association (EFMA) found that “11% had not started any activities in AI.”²¹ Six years later, a 2023 survey by Ernst & Young found that only 1% of financial services leaders reported that their organizations were not deploying AI in some manner, with 55% expressing optimism about their organization’s use of AI in the short term.²² Surges in use are extending into the corporate and retail banking sectors, according to

²⁰ Michelle Seng Ah Lee, Luciano Floridi, and Alexander Denev, “Innovating with Confidence: Embedding AI Governance and Fairness in a Financial Services Risk Management Framework,” Berkeley Technology Law Journal, 2020, https://btlj.org/wp-content/uploads/2020/01/Lee_Web.pdf; Jania Okwechime, “How Artificial Intelligence is Transforming the Financial Services Industry,” Deloitte Risk Advisory Insights, <https://www.deloitte.com/content/dam/assets-zone1/ng/en/docs/services/risk-advisory/2023/ng-how-artificial-intelligence-is-transforming-the-financial-services-industry.pdf>.

²¹ Lee et al. *supra* n. 20, citing Louise Brett et al., AI and You: Perceptions of Artificial Intelligence from the EMFA financial services industry, Deloitte, April 2017, <https://www2.deloitte.com/content/dam/Deloitte/cn/Documents/technology/deloitte-cn-tech-ai-and-you-en-170801.pdf>.

²² EY Americas, “EY Survey: AI adoption among financial services leaders universal, amid mixed signals of readiness,” December 11, 2023, https://www.ey.com/en_us/news/2023/12/ey-survey-ai-adoption-among-

2023 research conducted by McKinsey.²³ A Deloitte study on the use of AI by financial services firms also suggested its strategic positioning in core business units and extension to client engagement.²⁴

Meanwhile, banking firms are expected to earn between \$200 billion to \$340 billion in their adoption and use of generative AI.²⁵

Generally, AI models can support various use cases in the financial services industry, including fraud detection and prevention, customer relationship management, predictive analytics, and credit risk management.²⁶ More financial institutions, like trading and equity firms, are also leveraging AI tools, many of which are being integrated into functional tasks, including operations, risk management, compliance, and marketing.²⁷ The workplaces of the financial services industry are also shifting toward more automated functions; in the case of generative AI, a recent report by the MIT Technology Review suggests that the implementation of such technology has enabled “cutting costs by freeing employees from low-value, repetitive work.”²⁸

Increasing applications of AI will yield both potential benefits and risks to industry, individuals, and the broader public, which rely on the financial services marketplace. On the one hand, AI can aid in

[financial-services#:~:text=Nearly%20all%20\(99%25\)%20of,GenAI\)%20specifically%20within%20their%20organization.](#)

²³ Vishnu Kamalnath, Larry Lerner, Jared Moon, Gökan Sari, and Vik Sohoni, “Capturing the full value of generative AI in banking,” McKinsey & Company, December 5, 2023, <https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking>.

²⁴ “AI leaders in financial services,” Deloitte, August 13, 2019, <https://www.deloitte.com/global/en/our-thinking/insights/industry/financial-services/artificial-intelligence-ai-financial-services-frontrunners.html>.

²⁵ “Scaling gen AI in banking: Choosing the best operating model,” McKinsey & Company, March 22, 2024, <https://www.mckinsey.com/industries/financial-services/our-insights/scaling-gen-ai-in-banking-choosing-the-best-operating-model>.

²⁶ Okwechime, *supra* n. 20.

²⁷ Nvidia, “State of AI in Financial Services: 2024 Trends,” <https://resources.nvidia.com/en-us-2024-fsi-survey/ai-financial-services>.

²⁸ MIT Technology Review Insights, “Finding value in generative AI for financial services,” MIT Technology Review, November 26, 2023, <https://www.technologyreview.com/2023/11/26/1083841/finding-value-in-generative-ai-for-financial-services/>.

fraud detection, and, in some instances, make lending more accessible by enabling nontraditional scoring and credit assessments for those with insufficient credit history.²⁹ On the other hand, the use of AI in similarly situated, high-impact decision making scenarios could also translate into greater biases from autonomous decisions, which preference certain groups of applicants over others—creating differential treatment and potentially unlawful discrimination in the provision of financial services.³⁰ As with any other technology, AI can bolster efficiency, while simultaneously enabling the scale of fraudsters and scammers who may use the tools for ill will.³¹

How AI models are deployed throughout the financial services sector is highly relevant to the Commodity Futures Trading Commission (CFTC), whose mission is focused on promoting the integrity, resilience, and vibrancy of U.S. derivatives markets through sound regulation.³² In September 2022, CFTC Commissioner Goldsmith Romero reestablished the Technology Advisory Committee (TAC) to explore a range of advanced technology issues affecting the agency, including digital assets and blockchain technology, cybersecurity, and emerging and evolving technologies, such as AI.³³ Since its inception, the TAC has held hearings with experts from the White House, the U.S. Department of Commerce, and academia to report on governmental activities, responsible AI, and the latest research

²⁹ “Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services,” Virtual Hearing Before the Task Force on Artificial Intelligence of the Committee on Financial Services, U.S. House of Representatives, One Hundred Seventeenth Congress, First Session, May 7, 2021, Serial No. 117-23, <https://www.govinfo.gov/content/pkg/CHRG-117hrg44838/pdf/CHRG-117hrg44838.pdf>.

³⁰ Ibid.

³¹ Michael Stratford, “Banking regulator warns of ‘potential explosion’ of AI-fueled financial fraud,” Politico Pro, April 4, 2024, <https://subscriber.politicopro.com/article/2024/04/banking-regulator-warns-of-potential-explosion-of-ai-fueled-financial-fraud-00150576>.

³² CFTC Website, <https://www.cftc.gov/About/AboutTheCommission>.

³³ Renewal of the Technology Advisory Committee, 87 Fed Reg. 56003, September 13, 2022; Opening Statement of Commissioner Christy Goldsmith Romero at the Technology Advisory Committee on DeFi, Responsible Artificial Intelligence, Cloud Technology & Cyber Resilience, March 22, 2023, <https://www.cftc.gov/PressRoom/SpeechesTestimony/romerostatement032223>.

around the potential for super-manipulative chatbots in financial markets.³⁴ The TAC has also delivered a comprehensive report on decentralized finance (DeFi), which shared a series of recommendations to aid the agency in its oversight.³⁵ The Subcommittee has been tasked with the exploration of AI, and more advanced models like generative AI. A list of full members is included at the start of the report. For this report, for simplicity, we will collectively refer to “CFTC-registered entities” when discussing entities registered with or otherwise regulated by the Commission, including exchanges, clearinghouses, banks, and others acting as futures commission merchants, swap dealers, as well as managed funds and advisors, introducing brokers, and retail foreign exchange dealers.

To date, the CFTC has already issued a consumer advisory urging wariness of AI investing scams, reporting that fraudsters are making claims, often amplified by social media platforms, that AI-assisted investment tools can generate huge returns for investors.³⁶ The CFTC is also actively exploring the development of regulations and guidance for the use of AI in its regulated markets, recently releasing a request for comment on the definition of AI, its applications—such as trading, risk management, and compliance—and its risks—such as market manipulation, fraud, and bias.³⁷

As more financial organizations and adjacent industries, including third-party service providers, are leveraging AI across the sector, the collection and report out on AI use within institutional trading

³⁴ Ibid.; See also Opening Statement of Commissioner Christy Goldsmith Romero at the Technology Advisory Committee on Responsible AI in Financial Services, DAOs and Other DeFi & Cyber Resilience, July 18, 2023, <https://www.cftc.gov/PressRoom/SpeechesTestimony/romerostatement071823>; Opening Statement of Commissioner Christy Goldsmith Romero, Sponsor of the CFTC Technology Advisory Committee, on Responsible Artificial Intelligence, Cyber Resilience & Decentralized Finance, Jan. 8, 2024, <https://www.cftc.gov/PressRoom/SpeechesTestimony/romerostatement010824>.

³⁵ “Decentralized Finance,” Subcommittee on Digital Assets and Blockchain Technology, Technology Advisory Committee of the Commodity Futures Trading Commission, January 8th, 2024, https://www.cftc.gov/media/10106/TAC_DeFiReport010824/download.

³⁶ CFTC, “CFTC Customer Advisory Cautions the Public to Beware of Artificial Intelligence Scams,” Release Number 8854-24, January 25, 2024, <https://www.cftc.gov/PressRoom/PressReleases/8854-24>.

³⁷ CFTC, “CFTC Staff Releases Request for Comment on the Use of Artificial Intelligence in CFTC Regulated Markets,” January 25, 2024, Release Number 8853-24, <https://www.cftc.gov/PressRoom/PressReleases/8853-24>.

contexts in the United States are important and timely, especially since financial markets are of great importance to national security interests. The Subcommittee drafted this report with more granular focus on the opportunities and risks of increased AI applications and use cases among CFTC-registered entities. This Report finds that, without appropriate industry engagement and relevant guardrails (some of which have been outlined in existing national policies), potential vulnerabilities from using AI applications and tools within the CFTC, by regulated firms and exchanges, as well as third parties, could erode public trust in financial markets, services, and products. These factors can also make the supply chain less predictable.³⁸ The Subcommittee also opines on how the availability and use of more capable technologies, like generative AI, should also be on the radar of the CFTC as the agency serves to advance stakeholder engagement and offer potential guidance on AI applications with the most utility for their industry sectors.

The report describes various functions and operations of AI in known and anticipated institutional activities in the financial industry, and provides definitions, which are often associated with the technical and conceptual understandings of emerging and evolving technologies (see “Glossary”). The report presents a series of actual and hypothetical use cases, which may introduce novel risks to companies, markets, and investors, especially in high-impact, autonomous decision-making scenarios. The report further examines how humans are engaged in the design, interrogation, and oversight of autonomous models used in financial services, along with the possible biases in the training data and

³⁸ Unlike the Securities and Exchange Commission’s enabling statutes, the Commodity Exchange Act does not give the CFTC an explicit investor protection mandate, which creates a potential regulatory gap. With more retail customers entering the CFTC’s regulated markets, a new definition for retail customers may be needed to ensure customer protection. Keynote Remarks of Commissioner Christy Goldsmith Romero at the Futures Industry Association, Asia Derivatives Conference, Singapore, November 30, 2022, <https://www.cftc.gov/PressRoom/SpeechesTestimony/oparomero4>; Executive Order 14110, *supra* n. 2.

outputs that may arise. The report concludes with a series of recommendations that work toward more responsible AI systems with greater transparency, and oversight to appropriately safeguard financial markets.

There is still much more to be discovered and explored with the proliferation of AI among entities regulated by the CFTC. However, the Subcommittee is making the proposed recommendations to help provide guidance and a framework that can inform policy discussions and contribute to fostering safe, trustworthy, and responsible AI systems.

II. Defining AI and Responsible AI

Current definitions of AI vary, but there is some consensus around the core definition of the term referring to a science and a technology that aims to solve problems in a way that would be considered intelligent, as if it were a human being. Recent regulatory frameworks tend to lean on commonly referred definitions of AI systems. For example, in this definition by the Organisation for Economic Co-operation and Development (OECD): “[a]n AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.”³⁹ The definition of AI has also included many techniques: from rule-based applications to machine learning models to deep learning and generative AI, to name a few. There are also many interpretations of the terms relative to AI models, which are fleshed out in the Glossary.

Throughout the report, we also mention generative AI as a significant trendsetter in the future design, governance, and application of AI systems. It is worth pointing out its real-world, functional

³⁹ “OECD AI Principles Overview,” OECD.AI Policy Observatory, May 2019, <https://oecd.ai/en/ai-principles>.

applications for the CFTC in its decision-making. Because it is still early in the technology's deployment throughout the institutional side of financial services, the previous findings may have to be updated.

What is Responsible AI?

Adding to definitional challenges is use of terms that include "Responsible AI," "Trustworthy AI" (used by the European Commission), or "Ethical AI." These different terms are generally understood to refer to AI systems or applications that address issues and concerns with their use.

As referenced in this report, the definitions of "Responsible AI"⁴⁰ include five typical properties which speak to how AI models are designed and deployed:

- *Fairness* refers to processes and practices that ensure that AI does not make discriminatory decisions or recommendations.
- *Robustness* ensures that AI is not vulnerable to attacks to its performance.
- *Transparency* refers to sharing information that was collected during development and that describes how the AI system has been designed and built, and what tests have been done to check its performance and other properties.
- *Explainability* is the ability of the AI system to provide an explanation to users and other interested parties who inquire about what led to certain outputs in the AI's modelling. This is essential to generate trust in AI from users, auditors, and regulators.
- *Privacy* ensures that AI is developed and used in a way that protects users' personal information.

Over the years, AI techniques have significantly evolved, from rule-based systems where developers were explicitly coding the steps needed to solve problems, to machine learning systems where the AI

⁴⁰ What is responsible AI?, IBM, <https://www.ibm.com/topics/responsible-ai>.

system learns from extensive data about previously solved problems. These techniques significantly expanded the types of AI applications, which now provide the capability to interpret content such as images and text, and to generate predictions or classifications due to improved data learning processes.

What is generative AI?

Recent machine learning advances have led to the development of generative AI, which not only excels at interpreting content, but can also generate content (such as images, videos, and text). ChatGPT is a typical example of a system using generative AI, using a form of generative AI known as large language models (LLMs), where the AI model learns from large quantities of human language text how to generate a response to a textual prompt.⁴¹

Besides the issues and risks introduced earlier, new issues and concerns emerge with the use of generative AI, particularly if the generated content is deemed harmful (offensive, toxic, violent, obscene); the AI system “hallucinates” (that is, it can generate false content even if all training data is true); and/or humans use AI to produce fake content which is not distinguishable from reality, commonly referred to as “deepfakes.”

Generative AI also raises a host of legal implications for civil and criminal liability, including copyright infringement when the models are trained on copyrighted data and can generate and share part of its data.⁴² and misinformation when it generates deepfakes that can be used deceptively.⁴³ Many

⁴¹ “What are large language models?” IBM, <https://www.ibm.com/topics/large-language-models>.

⁴² Gil Appel, Juliana Neelbauer, and David A. Schweidel, “Generative AI Has an Intellectual Property Problem,” Harvard Business Review, April 7, 2023, <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>; Tate Ryan-Mosley, “How generative AI is boosting the spread of disinformation and propaganda,” MIT Technology Review, October 4, 2023, <https://www.technologyreview.com/2023/10/04/1080801/generative-ai-boosting-disinformation-and-propaganda-freedom-house/>.

⁴³ Danka Delić, “The deepfake technology behind a \$35 million bank heist in Hong Kong,” ProPrivacy, February 25, 2022, <https://proprivacy.com/privacy-news/deepfake-technology-used-in-hong-kong-bank-heist>.

content creators are currently fighting court battles over the use of their copyrighted materials in generative AI training and, in some cases, the alleged exact reproduction of their works by AI systems.⁴⁴ The plaintiffs in these cases include authors,⁴⁵ visual artists,⁴⁶ music publishers,⁴⁷ programmers,⁴⁸ and even major print media outlets, like newspapers.⁴⁹ The results of these lawsuits, which will have an enormous impact on the futures of AI development and copyright law, will hinge crucially on the degree to which the use of publicly available data (and in some instances, privately held information) constitutes fair use under settled law.⁵⁰ And much is still unknown on how the technology will impact the intellectual property of companies whose trade secrets may be exposed to emerging and evolving technologies.

In addition to questions of ownership, it is also unclear who, if anyone, is legally liable for misinformation generated by AI. Although existing laws protect online platforms, including AI developers, from liability for third-party content generated by users,⁵¹ there is currently a bill that

⁴⁴ Blake Brittain, “How copyright law could threaten the AI industry in 2024,” Reuters, January 2, 2024, <https://www.reuters.com/legal/litigation/how-copyright-law-could-threaten-ai-industry-2024-2024-01-02/>.

⁴⁵ Alexandra Alter and Elizabeth A. Harris, “Franzen, Grisham and Other Prominent Authors Sue OpenAI,” The New York Times, September 20, 2023, <https://www.nytimes.com/2023/09/20/books/authors-openai-lawsuit-chatgpt-copyright.html>.

⁴⁶ Jocelyn Noveck and Matt O’Brien, “Visual artists fight back against AI companies for repurposing their work,” Associated Press, August 31, 2023, <https://apnews.com/article/artists-ai-image-generators-stable-diffusion-midjourney-7ebcb6e6ddca3f165a3065c70ce85904>.

⁴⁷ Emilia David, “Universal Music sues AI company Anthropic for distributing song lyrics,” The Verge, October 19, 2023, <https://www.theverge.com/2023/10/19/23924100/universal-music-sue-anthropic-lyrics-copyright-katy-perry>.

⁴⁸ Thomas Claburn, “GitHub Copilot copyright case narrowed but not neutered,” The Register, January 12, 2024, https://www.theregister.com/2024/01/12/github_copilot_copyright_case_narrowed/.

⁴⁹ Michael M. Grynbaum and Ryan Mac, “The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work,” December 27, 2023, <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.

⁵⁰ Cala Coffman, “Does the Use of Copyrighted Works to Train AI Qualify as a Fair Use?,” Copyright Alliance, April 11, 2023, <https://copyrightalliance.org/copyrighted-works-training-ai-fair-use/>.

⁵¹ Barbara Ortutay, “What you should know about Section 230, the rule that shaped today’s internet,” PBS, February 21, 2023, <https://www.pbs.org/newshour/politics/what-you-should-know-about-section-230-the-rule->

proposes to remove these protections for generative AI chatbots.⁵² This legal debate is also beginning to play out in other countries. Recently, a generative AI-powered Air Canada bot misinformed a customer about the availability of a discount. When taken to the Canadian Civil Resolution Tribunal, the Tribunal rejected the argument that the AI was a “separate legal entity that is responsible for its own actions.”⁵³ Elsewhere, an Australian mayor considered but ultimately dropped a defamation suit against OpenAI, whose model incorrectly claimed he was jailed in a bribery scandal.⁵⁴ These and other examples will continue to arise in AI use cases. But this does not suggest that trading companies, and other similarly situated entities may be insulated from these and other effects, especially as AI models become more dual-purposed, and readily available for various use cases in financial services.

AI Governance

Such advancements in AI have also led to greater governance frameworks and practices. Governance can differ from direct regulation by setting general guiding principles for standards and practices; examples include the Office of Science and Technology Policy’s “Blueprint for an AI Bill of Rights,” which established five principles to inform AI governance,⁵⁵ and the NIST AI Risk Management

[that-shaped-todays-internet](#); Matt Schruers, “Myths and Facts about Section 230,” Disruptive Competition Project, October 16, 2019, <https://www.project-disco.org/competition/101619-myths-and-facts-about-section-230/>.

⁵² Sen. Josh Hawley, “A bill to waive immunity under section 230 of the Communications Act of 1934 for claims and charges related to generative artificial intelligence,” 118th Congress, June 14, 2023, <https://www.congress.gov/bill/118th-congress/senate-bill/1993/>.

⁵³ Marisa Garcia, “What Air Canada Lost In ‘Remarkable’ Lying AI Chatbot Case,” Forbes, February 19, 2024, <https://www.forbes.com/sites/marisagarcia/2024/02/19/what-air-canada-lost-in-remarkable-lying-ai-chatbot-case/?sh=2bef14e0696f>.

⁵⁴ David Swan, “Australian mayor abandons world-first ChatGPT lawsuit,” The Sydney Morning Herald, February 12, 2024, <https://www.smh.com.au/technology/australian-mayor-abandons-world-first-chatgpt-lawsuit-20240209-p5f3nf.html>.

⁵⁵ White House Office of Science and Technology Policy, “Blueprint for an AI Bill of Rights,” <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

Framework.⁵⁶ That is, the typical properties and deployment practices of AI are beginning to require governance structures, which enable relevant guardrails that protect both the consumers and contexts in which the technology is deployed.

Governance also applies to various stages of the technology, which are worth pointing out. The *first* type of AI governance is focused on the lifecycle of the model, which includes data collection, data labelling, model training, model testing, passing from an AI model to an AI system, system deployment, and system monitoring.⁵⁷ At every phase of this lifecycle, there are a series of checks, consultations, reporting, and testing steps that need to be conducted to make sure that the resulting AI is trustworthy and responsible. Areas of focus for whole-lifecycle AI governance should likely include the issues of “value alignment” and “humans in the loop.”. Value alignment refers to restriction of an AI model so that it can only pursue goals that follow values aligned to human values. Current research (including from members of the broader TAC) in training AI systems to operate ethically offers a potential technological means to prevent alignment failure between metrics and parameters of AI model developers and the risk metrics and parameters of the business units adopting the AI model.⁵⁸ It is crucial that the Commission avail itself of computer science research to distinguish value alignment from safety requirements to comprehensively assess risk in AI systems.

⁵⁶ National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” NIST AI 100-1, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

⁵⁷ Andrew Ferlitsch, “Making the machine: the machine learning lifecycle,” Google Cloud Blog, February 6th, 2019, <https://cloud.google.com/blog/products/ai-machine-learning/making-the-machine-the-machine-learning-lifecycle>.

⁵⁸ Ritesh Noothigattu, Djallel Bouneffouf, Nicholas Mattei, Rachita Chandra, Piyush Madan, Kush R. Varshney, Murray Campbell, Mohinder Singh, and Francesca Rossi, “Teaching AI Agents Ethical Values Using Reinforcement Learning and Policy Orchestration,” IBM Journal of Research and Development, 2019, p. 1, https://www.cs.cmu.edu/~rnoothig/papers/policy_orchestration.pdf.

Second, AI governance can refer to the corporate or public policies and regulations on the use of AI, especially in regulated industries, like financial services, employment, and health care. In this latter definition of governance, these policies and/or regulations are tied to the concept of responsible AI, since their aim is usually to make sure that only responsible AI is adopted in the market, especially in high-risk scenarios that include, for example, human resources, the health system, and the financial domain. For example, the AI Act adopted by the European Union in March 2024 is a regulation based on the concept of trustworthy AI, as defined by the European Commission High Level Expert Group on AI.⁵⁹

Third, AI governance is often what companies are putting in place internally. At this point, most companies creating AI systems have already started to define and adopt AI governance frameworks that include risk case assessment processes, internal governance bodies, engagement activities, software tools, testing methods, products, and platforms. Increasingly, AI governance is considered a necessary part of building and using AI systems.

Governments and international regulatory bodies can set guardrails for use of emerging technologies and can influence or set requirements for AI governance. The next section describes the U.S. legislative and regulatory developments related to AI generally, as well as international developments, before discussing rules and regulations applicable to use of AI in financial services specifically.

III. The AI Policy Landscapes – U.S. and Abroad

Domestic Policies

In recent years, the United States has developed foundational proposals, which have been supported through voluntary commitments with the private sector, Congressional hearings, and cross-

⁵⁹ “Trustworthy AI,” European Commission AI Watch, https://ai-watch.ec.europa.eu/topics/trustworthy-ai_en.

sector roundtable discussions, among other things, regarding domestic AI governance and expected use cases of AI by federal agencies.⁶⁰ In October 2022, the White House Office of Science and Technology Policy (OSTP) published the Blueprint for an AI Bill of Rights (Blueprint).⁶¹ In the Blueprint, OSTP outlined nonbinding guidelines for consumer protections regarding AI. OSTP categorized these guidelines in accordance with five overarching principles: safe and effective systems; algorithmic discrimination protections; data privacy; notice and explanation; and human alternatives, consideration, and fallback.

Compliance with the Blueprint’s guidelines is voluntary, but the White House announced that it would be followed by related actions from multiple government agencies.⁶² At the time of the Blueprint’s release, at least a dozen agencies had issued some sort of binding guidance for the use of automated systems in the industries under their jurisdiction, but the detail and quality of these policies varied substantially.⁶³

The lack of mandatory, enforceable guidelines, as well as the Blueprint’s carve-out of law enforcement and national security—two sectors in which the harms of AI can be particularly severe—limit the direct impact of this guidance.⁶⁴

⁶⁰ See CFTC, “Commissioner Goldsmith Romero Announces Agenda for January 8 Technology Advisory Committee Meeting on Artificial Intelligence, Cybersecurity, Decentralized Finance,” Release Number 8846-24, January 5, 2024, <https://www.cftc.gov/PressRoom/PressReleases/8846-24>.

⁶¹ Blueprint, *supra* n. 55.

⁶² Kay Firth-Butterfield, Karen Silverman, and Benjamin Larsen, “Understanding the US ‘AI Bill of Rights’ - and how it can help keep AI accountable,” World Economic Forum, October 14, 2022, <https://www.weforum.org/agenda/2022/10/understanding-the-ai-bill-of-rights-protection/>.

⁶³ Nicol Turner Lee and Jack Malamud, “Opportunities and blind spots in the White House’s blueprint for an AI Bill of Rights,” The Brookings Institution, December 19, 2022, <https://www.brookings.edu/articles/opportunities-and-blind-spots-in-the-white-houses-blueprint-for-an-ai-bill-of-rights/>.

⁶⁴ *Ibid.*

In January 2023, NIST published its Congressionally mandated AI Risk Management Framework (RMF).⁶⁵ Similar to OSTP’s Blueprint, NIST’s AI RMF provides a voluntary roadmap toward identifying and mitigating risks of AI in concrete contexts and sets forth seven characteristics of trustworthy AI: safe; secure and resilient; explainable and interpretable; privacy-enhanced; fair; accountable and transparent; and valid and reliable.⁶⁶

One year after the release of the Blueprint for an AI Bill of Rights, on October 30, 2023, President Biden issued his an “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.”⁶⁷ This sweeping action implemented many of the principles included in the Blueprint, issuing guidance for agencies’ use and procurement of AI systems, requiring that developers of certain AI systems share safety test results with the government, instructing NIST to set standards for adversarial testing, and providing “clear guidance to landlords, Federal benefits programs, and federal contractors to keep AI algorithms from being used to exacerbate discrimination.”⁶⁸ The order also requires cloud service providers to report foreign customers to the government and recommends the use of digital watermarks to identify AI-generated media.⁶⁹

To ensure that federal agencies are working towards more responsible and trustworthy AI, the Office of Management and Budget (OMB) released a draft of guidance for the federal agencies required

⁶⁵ National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” NIST AI 100-1, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

⁶⁶ Cameron F. Kerry, “NIST’s AI Risk Management Framework plants a flag in the AI debate,” The Brookings Institution, February 15, 2023, <https://www.brookings.edu/articles/nists-ai-risk-management-framework-plants-a-flag-in-the-ai-debate/>.

⁶⁷ Executive Order 14110, *supra* n. 2.

⁶⁸ The White House, “FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence,” October 30, 2023, <https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/>.

⁶⁹ Cecilia Kang and David E. Sanger, “Biden Issues Executive Order to Create A.I. Safeguards,” The New York Times, October 30, 2023, <https://www.nytimes.com/2023/10/30/us/politics/biden-ai-regulation.html>.

to implement the Order’s provisions.⁷⁰ This action served as an addendum to the White House EO. The proposed guidance—which is still in draft form—would require federal agencies to develop and publish individual strategies for the use of AI, implement impact assessments for AI systems, and designate a Chief AI Officer for each agency to manage AI oversight. Senators Jerry Moran (R-KS) and Mark Warner (D-VA) have since proposed the Financial Artificial Intelligence Risk Reduction Act (FAIRR), which would codify some of the provisions of President Biden’s Executive Order and translate the NIST AI RMF into mandatory commitments that federal agencies would be obliged to use in procurement of AI systems.⁷¹ Representative Ted Lieu (D-CA) has also introduced a companion bill to the FAIRR Act in the House of Representatives. More recently the Chairman of the House Committee on Oversight and Accountability, Representative James Comer (R-KY), and Ranking Member Jamie Raskin (D-MD) proposed the Federal AI Governance and Transparency Act, which would codify AI acquisition and oversight safeguards, establish AI governance charters for federal agencies, and create a notification process for individuals who are affected by autonomous decision-making processes.⁷²

⁷⁰ The White House, “OMB Releases Implementation Guidance Following President Biden’s Executive Order on Artificial Intelligence,” November 1, 2023, <https://www.whitehouse.gov/omb/briefing-room/2023/11/01/omb-releases-implementation-guidance-following-president-bidens-executive-order-on-artificial-intelligence/>; Shalanda D. Young, Proposed Memorandum for the Heads of Executive Departments and Agencies: “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Office of Management and Budget, <https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf>.

⁷¹ Rebecca Kern and Brendan Bordelon, “Senators push to give Biden’s AI order more teeth,” Politico, November 2, 2023, <https://www.politico.com/news/2023/11/02/senate-ai-bill-biden-executive-order-00124893>; FAIRR Act, S.3554, 118th Congress (2023), <https://www.congress.gov/bill/118th-congress/senate-bill/3554?q=%7B%22search%22%3A%22FAIRR%22%7D&s=2&r=1>.

⁷² Federal Artificial Intelligence Risk Management Act, H.R.6936, 118th Congress (2024), <https://www.congress.gov/bill/118th-congress/house-bill/6936/text>; Federal AI Governance and Transparency Act, H.R.7532, 118th Congress (2024), <https://www.congress.gov/bill/118th-congress/house-bill/7532/text>; Brennan Center for Justice, “Artificial Intelligence Legislation Tracker,” <https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker>.

In addition to federal actions, there have also been flurries of state efforts to govern the technology. From January to October 2023, 24 states⁷³ introduced a total of 190 AI-related bills, which is a 440% increase compared to 2022 and more than the previous two years combined.⁷⁴ These bills covered diverse topics ranging from algorithmic bias (New York S-7623⁷⁵) and personal privacy (Massachusetts S-227⁷⁶) to state AI Task Forces (Illinois HB-3563⁷⁷), Offices (Connecticut SB-1103⁷⁸), and advisory councils (Texas HB-2060⁷⁹) and licensing requirements for advanced AI systems (New York A-8195⁸⁰). Naturally, a particular focus was on newly popular generative AI technology, with bills addressing election interference (Michigan HB-5144,⁸¹ Wisconsin SB-644,⁸² New York A-6790⁸³),

⁷³ “US State-by-State AI Legislation Snapshot,” BCLP, February 12, 2024, <https://www.bclplaw.com/en-US/events-insights-news/2023-state-by-state-artificial-intelligence-legislation-snapshot.html>.

⁷⁴ “BSA Analysis: State AI Legislation Surges by 440% in 2023,” BSA | The Software Alliance, September 27, 2023, <https://www.bsa.org/news-events/news/bsa-analysis-state-ai-legislation-surges-by-440-in-2023>.

⁷⁵ Relates to restricting the use of electronic monitoring and automated employment decision tools, 2023 (NY S7623A), <https://www.nysenate.gov/legislation/bills/2023/S7623/amendment/A>.

⁷⁶ An Act establishing the Massachusetts Information Privacy and Security Act, 2023 (MA S.227), <https://malegislature.gov/Bills/193/S227>.

⁷⁷ DOIT-AI TASK FORCE, 2023 (IL HB3563), <https://www.ilga.gov/legislation/BillStatus.asp?DocNum=3563&GAID=17&DocTypeID=HB&SessionID=112&GA=103>.

⁷⁸ An Act concerning artificial intelligence, automated decision-making, and personal data privacy, 2023 (CT Senate Bill No. 1103), <https://www.cga.ct.gov/2023/cbs/S/pdf/SB-1103.pdf>.

⁷⁹ Relating to the creation of the artificial intelligence advisory council, 2023 (TX House Bill 2060), <https://capitol.texas.gov/BillLookup/History.aspx?LegSess=88R&Bill=HB2060>.

⁸⁰ Relates to enacting the “advanced artificial intelligence licensing act,” 2023 (NY A8195), <https://www.nysenate.gov/legislation/bills/2023/A8195>.

⁸¹ Public Act 265 of 2023, 2023 (MI House Bill 5144), [http://www.legislature.mi.gov/\(S\(guo5a4ufjqycm2nxog42pw2v\)\)/mileg.aspx?page=GetObject&objectname=2023-HB-5144](http://www.legislature.mi.gov/(S(guo5a4ufjqycm2nxog42pw2v))/mileg.aspx?page=GetObject&objectname=2023-HB-5144).

⁸² An Act to amend 11.1303, 2024 (WI Senate Bill 644), <https://docs.legis.wisconsin.gov/2023/proposals/reg/sen/bill/sb644>.

⁸³ Prohibits the creation and dissemination of synthetic media within sixty days of an election with intent to unduly influence the outcome of an election, 2023 (NY State Assembly Bill A6790A), <https://www.nysenate.gov/legislation/bills/2023/A6790/amendment/A>.

deepfakes (New Jersey A-5512⁸⁴), notices of AI utilization by state agencies (California SB-313⁸⁵) and labeling of AI-generated content in advertisements (New York A-216⁸⁶), and broader regulation of large generative AI models (Massachusetts S-31⁸⁷).

Actions taken by the state of California are of particular interest, which, as home to 32 of Forbes' top 50 global AI companies,⁸⁸ may be uniquely positioned to have an outsized impact on AI governance. In September 2023, Governor Gavin Newsom issued his own Executive Order which directed state agencies to, on precise timelines, examine risks, reform public sector procurement for AI, consider the privacy and fairness impacts of state AI use, and analyze potential effects on the workforce.⁸⁹ More recently, in early February 2024, Senator Scott Wiener (D-San Francisco) introduced a bill (CA SB-1047⁹⁰) that proposes to require pre-deployment safety testing of future frontier models, establish a public computational resource to empower researchers, small startups, and community

⁸⁴ Establishes Deep Fake Technology Unit in DLPS, 2023 (NJ Bill A5512), <https://www.njleg.state.nj.us/bill-search/2022/A5512>.

⁸⁵ An act to add Chapter 5.9 (commencing with Section 11549.80) to Part 1 of Division 3 of Title 2 of, the Government Code, relating to state government, 2023 (CA Senate Bill-313), https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB313.

⁸⁶ Relates to requiring advertisements to disclose the use of synthetic media, 2023 (NY State Assembly Bill A216B), <https://www.nysenate.gov/legislation/bills/2023/A216/amendment/B>.

⁸⁷ An Act drafted with the help of ChatGPT to regulate generative artificial intelligence models like ChatGPT, 2023 (MA Bill S.31), <https://malegislature.gov/Bills/193/S31>.

⁸⁸ Kenrick Cai, "AI 50," Forbes, April 11, 2023, <https://www.forbes.com/lists/ai50/>.

⁸⁹ Darrell M. West, "California charts the future of AI," Brookings, September 12, 2023, <https://www.brookings.edu/articles/california-charts-the-future-of-ai/>; Office of Governor Gavin Newsom, "Governor Newsom Signs Executive Order to Prepare California for the Progress of Artificial Intelligence," September 6, 2023, <https://www.gov.ca.gov/2023/09/06/governor-newsom-signs-executive-order-to-prepare-california-for-the-progress-of-artificial-intelligence/>; State of California Executive Department, Executive Order N-12-23, September 6, 2023, <https://www.gov.ca.gov/wp-content/uploads/2023/09/AI-EO-No.12--GGN-Signed.pdf>.

⁹⁰ Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, 2024 (CA Senate Bill-1047).

groups to develop large AI models, and make developers answerable to the state Attorney General for severe harms due to irresponsible behavior and imminent public safety threats due to negligence.⁹¹

As federal government bodies explore their agendas for establishing the appropriate levels of legislation and guidance on AI, states will most likely continue to develop their own standards, and norms, which may potentially create a patchwork of incompatible public approaches to the technology's use cases.

More recent OMB Guidance

On March 28, 2024, the OMB released a new policy for federal agencies entitled "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," continuing its implementation of President Biden's AI Executive Order.⁹² The new guidance requires federal agencies to implement safeguards for AI use cases that impact "American's rights or safety" by December 1, 2024, and cease using any AI systems that prevent the application of such safeguards, "unless agency leadership justifies why doing so would increase risks to safety or rights overall or would create an unacceptable impediment to critical agency operations."⁹³ The guidance also requires federal agencies to release inventories of AI use cases, encourages agencies to "responsibly experiment with generative

⁹¹ "Senator Wiener Introduces Legislation to Ensure Safe Development of Large-Scale Artificial Intelligence Systems and Support AI Innovation in California," Office of Scott Wiener, February 8, 2024, <https://sd11.senate.ca.gov/news/20240208-senator-wiener-introduces-legislation-ensure-safe-development-large-scale-artificial>.

⁹² Shalanda D. Young, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," M-24-10 MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES, March 28, 2024, <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf>.

⁹³ The White House, "FACT SHEET: Vice President Harris Announces OMB Policy to Advance Governance, Innovation, and Risk Management in Federal Agencies' Use of Artificial Intelligence," March 28, 2024, <https://www.whitehouse.gov/briefing-room/statements-releases/2024/03/28/fact-sheet-vice-president-harris-announces-omb-policy-to-advance-governance-innovation-and-risk-management-in-federal-agencies-use-of-artificial-intelligence/>.

AI,” and directs agencies to designate Chief AI Officers and establish AI Governance Boards.⁹⁴ The OMB’s implementation work will continue with a request for information (RFI) on AI procurement.⁹⁵ As these national actions roll out, it will be imperative for the CFTC to take note, and respond accordingly in response to voluntary and mandatory guidance.

Global AI Policies

Despite the increasing policy activities in the U.S., for years, the European Union (EU) has been deeply focused on regulating emerging and evolving technologies, or some parts of the surrounding ecosystem including online data privacy. In 2016, the EU adopted the General Data Protection Regulation (GDPR), which Member States fully implemented by May 2018.⁹⁶ The GDPR primarily governs how businesses can handle individuals’ personal data, such as their names, location data, IP address, political opinions, and health information. The GDPR lays out seven key principles: lawfulness, fairness, and transparency; purpose limitation; data minimization; accuracy; storage limitation; integrity and confidentiality; and accountability.⁹⁷ Although the GDPR primarily concerns privacy, it has several direct and indirect implications for the use of AI systems as well. For instance, the GDPR requires systems to have a “human-in-the-loop”—meaning that a human overseer takes an active role in influencing, verifying, and approving the output of an AI system⁹⁸—for any automated decision with “legal effects or

⁹⁴ Ibid.

⁹⁵ Ibid.

⁹⁶ European Data Protection Supervisor, “The History of the General Data Protection Regulation,” https://edps.europa.eu/data-protection/data-protection/legislation/history-general-data-protection-regulation_en.

⁹⁷ Matt Burgess, “What is GDPR? The summary guide to GDPR compliance in the UK,” Wired, March 24, 2020, <https://www.wired.co.uk/article/what-is-gdpr-uk-eu-legislation-compliance-summary-fines-2018>.

⁹⁸ Cem Dilmegani, “Human-in-the-Loop Automation in 2024: Benefits & Applications,” AI Multiple Research, January 12, 2024, <https://research.aimultiple.com/human-in-the-loop-automation>.

similarly significant effects” and requires impact assessments for high-risk uses of data.⁹⁹ In 2020, the European Parliamentary Research Service conducted a study on the impact of the GDPR on AI, finding no incompatibility between the GDPR and AI development and noting that the GDPR’s guidance for data protection (and data protection impact assessments) are highly relevant for the development of AI systems.¹⁰⁰

In December 2023, the EU finalized its negotiations on the contents of the AI Act, which will be one of the first comprehensive AI legislative acts in the world (second only to regulations implemented by China in August).¹⁰¹ The Act takes a risk-based approach—scaling the rigor of its regulations according to the risk posed by a given application of AI—and bans both predictive policing and biometric systems that identify people using sensitive characteristics such as sexual orientation and race, and the indiscriminate scraping of faces from the internet, although it allows the use of “biometric identification” systems for some serious crimes.¹⁰² The AI Act also includes transparency requirements for “general purpose AI models” and steep fines for companies that fail to comply.¹⁰³

⁹⁹ Cameron F. Kerry, “Protecting privacy in an AI-driven world,” Brookings, February 10, 2020, <https://www.brookings.edu/articles/protecting-privacy-in-an-ai-driven-world/>; “Data Protection Impact Assessment (DPIA),” GDPR, <https://gdpr.eu/data-protection-impact-assessment-template/>.

¹⁰⁰ European Parliamentary Research Service, “The impact of the General Data Protection Regulation (GDPR) on artificial intelligence,” Panel for the Future of Science and Technology, PE 641.530, June 2020, [https://www.europarl.europa.eu/RegData/etudes/STUD/2020/641530/EPRS_STU\(2020\)641530_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2020/641530/EPRS_STU(2020)641530_EN.pdf).

¹⁰¹ Morgan Meaker, “The EU Just Passed Sweeping New Rules to Regulate AI,” Wired, December 8, 2023, <https://www.wired.com/story/eu-ai-act/#:~:text=The%20European%20Union%20agreed%20on,racing%20to%20develop%20AI%20systems.>

¹⁰² Ibid; Kelvin Chan, “Europe agreed on world-leading AI rules. How do they work and will they affect people everywhere?” The Associated Press, December 11, 2023, <https://apnews.com/article/ai-act-artificial-intelligence-regulation-europe-06ab334caa97778770f5f57f4d904447>.

¹⁰³ Ibid.

Policy References to Financial Services Protections

As the U.S. and EU develop more grand public policies to structure how AI should be governed, other policies are also resonating among specific federal agencies. For example, the White House EO instructed the Consumer Financial Protection Bureau (CFPB) and the Federal Housing Finance Agency (FHFA) to require regulated entities to evaluate AI models used for collateral valuation and appraisal for bias against protected groups.¹⁰⁴ The CFPB also issued guidance on lenders using artificial intelligence to deny credit, explaining that lenders must use “specific and accurate reasons when taking adverse actions against consumers,” and the Securities and Exchange Commission (SEC) proposed new rules to address potential conflicts of interest posed using AI systems by brokers and investment advisers.¹⁰⁵ It is important to note that there is a strong nexus between the SEC and the CFTC’s respective remits and that many CFTC registrants are dually registered with the CFTC and the SEC. As a result, many entities regulated by the CFTC, if engaged in SEC-regulated activities, will need to comply with any SEC rulemaking on AI investing and related conflicts of interest.

It is important to note that the CFTC has for some time been researching AI’s evolution, taxonomies, applications, and governance. In 2019, CFTC staff published a primer on AI in financial markets that presented a brief history of the evolution of AI in computer science and its adoption in financial services. The primer also included a summary of the rise of Big Data that made AI possible and

¹⁰⁴ Joseph Kamyar, David Simon, Simon Toms, Eve-Christie Vermynck, “AI Insights: How Regulators Worldwide Are Addressing the Adoption of AI in Financial Services,” JDSupra, December 13, 2023, <https://www.jdsupra.com/legalnews/ai-insights-how-regulators-worldwide-1616649/>.

¹⁰⁵ Ibid; CFPB, “CFPB Issues Guidance on Credit Denials by Lenders Using Artificial Intelligence,” September 19, 2023, <https://www.consumerfinance.gov/about-us/newsroom/cfpb-issues-guidance-on-credit-denials-by-lenders-using-artificial-intelligence/>; Ken Kumayama, Stuart Levi, Anna Rips, and Resa Schlossberg, “SEC Proposes New Conflicts of Interest Rule for Use of AI by Broker-Dealers and Investment Advisers,” JDSupra, August 11, 2023, <https://www.jdsupra.com/legalnews/sec-proposes-new-conflicts-of-interest-8410102/>.

an AI case study.¹⁰⁶ These prior actions with the newly established TAC will be helpful in their pursuit of increased awareness by their internal and external stakeholders.

IV. Financial Services AI Use Cases

As suggested, AI has been used in impactful ways in the financial industry for more than two decades, although basic computational models and statistical methods, such as standard deviations and Bayesian regressions, have been in use since the 1980s to generate trading signals. Machine learning (ML) gained significant traction in the 2000s and early 2010s in high-frequency trading and risk modeling. This increase in traction was driven by increasing accessibility to data, dropping costs for data storage, and increasing computational power. As technology evolved, financial institutions, including banks, employed increasingly complex AI systems, such as neural networks powered by extensive credit card data to determine eligibility.

The use of AI quickly spread to other areas of finance, such as fraud detection, consumer credit scoring, and customer service. By mid-2010s, the impact of AI was transforming personal finance with the availability of robo-advisory platforms through startups and established financial institutions alike.¹⁰⁷ In addition, advancements in deep learning and reinforcement learning models continued to improve the financial industry's efficiency. With the increasing accessibility of foundation models, they are now able to power generative AI and revolutionize the financial industry. Generative AI can enable financial companies to unlock value by creating new products by analyzing vast amounts of previously untapped

¹⁰⁶ See LabCFTC, CFTC, A Primer on Artificial Intelligence in Financial Markets, p. 13,

https://www.cftc.gov/media/2846/LabCFTC_PrimerArtificialIntelligence102119/download.

¹⁰⁷ Brooke Southall, "What exactly are robo-advisors, why did they steal the 2014 show and what will a 2015 repeat take?," RIABiz, January 2, 2015, <https://riabiz.com/a/2015/1/2/what-exactly-are-robo-advisors-why-did-they-steal-the-2014-show-and-what-will-a-2015-repeat-take>.

unstructured textual data, which according to IDC's white paper "Untapped Value: What Every Executive Needs to Know About Unstructured Data," is estimated to comprise 80–90% of all existing data.¹⁰⁸ However, the unstructured data may not necessarily be terra incognita. That is, the use of such data requires effort to establish its provenance and acquire legal permission to use it. If these data controls are ignored, AI could replicate the "move fast and break things" practices of social media both in developing internal AI models and third-party models.

The combination of AI, programmable digital assets, and smart contracts may be potentially beneficial and could possibly create a financial system that can efficiently run complex tasks and enforce financial agreements without, or with limited, human intervention. For example, AI algorithms can trigger smart contracts to buy and sell assets when market conditions are met or freeze digital assets from further transfers when fraudulent activities are detected. Smart contracts can automatically record each step of an AI algorithm, providing a transparent and immutable audit trail for compliance or further training of AI algorithms.

As AI continues to learn from the trusted dataset, it can, in turn, adapt and optimize its algorithms as well as smart contracts to new market conditions. Additionally, digital assets on blockchain can also protect against fake digital assets, which are increasingly easy to create with generative AI. Users of digital assets can validate the authenticity of a digital asset by checking the issuing contract against a shared digital asset registry controlled by authorized parties. As automation and digitization proliferate in financial markets, it is crucial that markets simultaneously prioritize operational resilience, such as cybersecurity measures that are robust to AI-empowered cyberattacks. AI

¹⁰⁸ Holly Muscolino and Amy Machado, "Untapped Value: What Every Executive Needs to Know About Unstructured Data," IDC, August 2023, <https://images.g2crowd.com/uploads/attachment/file/1350731/IDC-Unstructured-Data-White-Paper.pdf>.

can monitor transactional data in real-time, identifying and flagging any unusual activities.¹⁰⁹ Advanced ML algorithms can aid in the prediction of future attack vectors based on existing patterns,¹¹⁰ providing an additional layer of cybersecurity. As foundation models get trained on the trusted digital asset transactions and smart contract code on blockchains, generative AI may be able to help developers write more secure code.

AI systems can be less beneficial and, in some instances, more dangerous if the potential challenges and embedded biases in AI models are regressive to financial gains, and in worst case scenarios, prompt significant market instabilities due to the interjection of misguided training data, or the goals of bad actors to disrupt markets. In these and other instances, having more robust voluntary guidance and more prescriptive public policies can raise assurances around AI safety among CFTC-regulated industries.

Other Considerations

For AI systems to be wholly effective, other important considerations include the quality of the data used to train financial service models, including the role of a human-in-the-loop to help identify and mitigate risks, and the potential for data and/or predictive biases, which could upend existing assumptions about financial markets. The Subcommittee also makes the case that in fully autonomous models, there may be no need for human engagement, or at least at the level of monitoring external biases. Finally, there will be an increasing need for skilled talent within the Commission, who can bring

¹⁰⁹ Benjamin Chou, "The Digital Sentry: How AI Will Revolutionize Financial Fraud Investigation," *Forbes*, July 3, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/07/03/the-digital-sentry-how-ai-will-revolutionize-financial-fraud-investigation/?sh=390279657e24>.

¹¹⁰ Arash Negahdari Kia, Finbarr Murphy, Barry Sheehan, and Darren Shannon, "A cyber risk prediction model using common vulnerabilities and exposures," *Expert Systems with Applications* 237 (March 2024): 121599, <https://doi.org/10.1016/j.eswa.2023.121599>.

expertise in the governance and rulemaking processes and represent the agency at AI-specific convenings.

Training data

Training data is the raw material from which a machine learning system learns patterns to make predictions. As such, ensuring high-quality training data is the most important part of the development process.¹¹¹ Certainly, other factors are necessary in designing AI, such as choosing the right type of algorithm¹¹² and ensuring that the system generalizes well outside the data it was trained on.¹¹³ However, if either the individual training samples or the broader dataset have problems, it does not matter which algorithm is chosen or which measures are taken to ensure good generalization; the raw learning material will be tainted, and consequently, the system will struggle to make accurate predictions upon deployment.

A Deloitte report on managing AI risks in investment management details four guidelines for ensuring data quality.¹¹⁴ First, they recommend automated quality assurance metrics that execute as the data are acquired and comb the dataset for missing values, outliers, and changes in volatility patterns. Second, they advocate ongoing monitoring processes that can regularly monitor, track, and report data failures and errors. Third, they recommend data remediation techniques to correct data errors and bring the quality closer to that of a known high-quality dataset. Finally, they propose quality

¹¹¹ “The Essential Guide to Quality Training Data for Machine Learning,” cloudfactory, 2020,

<https://www.cloudfactory.com/training-data-guide>.

¹¹² Hui Li, “Which machine learning algorithm should I use?,” The SAS Data Science Blog, December 9, 2020,

<https://blogs.sas.com/content/subconsciousmusings/2020/12/09/machine-learning-algorithm-use/>.

¹¹³ “Overfitting in Machine Learning: What It Is and How to Prevent It,” Elite Data Science, July 6, 2022,

<https://elitedatascience.com/overfitting-in-machine-learning>.

¹¹⁴ “Managing model and AI risks in the investment management sector,” Deloitte, 2023,

<https://www2.deloitte.com/content/dam/Deloitte/us/Documents/Advisory/us-advisory-deloitte-managing-model-and-ai-risks-in-the-investment-management-sector-july-2023.pdf>.

reporting and data signoff to monitor data quality at every stage in the acquisition and processing of the training data.

In addition to the quality of the data, the type of data itself also matters. The noisy nature of one-dimensional market prices can make accurate predictions more difficult. However, researchers have found that they can improve predictive accuracy by incorporating additional sources of information in their training data such as economic news, social media, and even weather information.¹¹⁵

It is also extremely important to ensure that the training data is representative of what the model will see upon deployment.¹¹⁶ If it is not, then the model may have trouble when confronted with patterns that are less common in the training data but more common in the real world. For example, facial recognition algorithms trained mostly on white faces have higher error rates on non-White faces, while similar algorithms trained mostly on Asian faces have higher error rates on non-Asian faces.¹¹⁷ Where the focus on biased data by certain demographics has been critical to the debates on algorithmic fairness, some aspects of this principled approach could apply to organizations' trading data.

If data quality and type are not prioritized during the training process, mistakes may occur. However, quality training data does not guarantee that an AI system may be less likely to err. Modern AI systems constructed by top companies with vast data access can and do make costly public errors¹¹⁸—

¹¹⁵ Shuo Sun, Rundong Wang, and Bo An, "Reinforcement Learning for Quantitative Trading," *ACM Transactions on Intelligent Systems and Technology* 14, 3 (2023) p. 12, <https://arxiv.org/pdf/2109.13851.pdf>.

¹¹⁶ Marvin Zhang, "Adaptation Based Approaches to Distribution Shift Problems," University of California, Berkeley, December 17, 2021, p. 17, <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2021/EECS-2021-262.pdf>.

¹¹⁷ National Academies of Sciences, Engineering, and Medicine, "Facial Recognition Technology: Current Capabilities, Future Prospects, and Governance," Washington, DC: The National Academies Press, 2024, <https://doi.org/10.17226/27397>.

¹¹⁸ Zhang, *supra* n. 116.

such as Alphabet’s \$90B market loss after its large multimodal model Gemini generated controversially inaccurate images¹¹⁹—many of which could have been prevented with human involvement.

The placement of humans

Where humans fit in the design, deployment, and oversight of models must be further considered when dealing with fully autonomous, financial systems. In particular, the role of a human-in-the-loop (HITL), and human-out-the-loop (HOTL) are all part of the process and will impact governance strategies.

- A. HITL: Putting humans in the loop can not only prevent costly mistakes, but also create a human-machine system that can be more capable than either of its parts¹²⁰ and mitigate data quality issues (if the algorithm learns from deployment experience).¹²¹ However, it has been suggested HITL is not always feasible or desirable. For example, it may be impossible for humans to be in the loop for high-frequency trading algorithms, which take advantage of market discrepancies on milli- to microsecond (one millionth of a second) timescales.¹²² It may also be less feasible for humans to be in the loop if the algorithm has a large output volume; for example, stopping the Gemini controversy would not have been possible solely with runtime oversight, due to the number of people that would be required to check every output. Entities must balance business opportunities and the risk of those opportunities,

¹¹⁹ Derek Saul, “Google’s Gemini Headaches Spur \$90 Billion Selloff,” February 26, 2024, <https://www.forbes.com/sites/dereksaul/2024/02/26/googles-gemini-headaches-spur-90-billion-selloff/?sh=387587fb72e4>.

¹²⁰ Samuel R. Bowman et al., “Measuring Progress on Scalable Oversight for Large Language Models,” arXiv, November 11, 2022, <https://arxiv.org/abs/2211.03540>.

¹²¹ Cem Dilmegani, “Human-in-the-Loop Automation in 2024: Benefits & Applications,” AIMultiple Research, January 12, 2024, <https://research.aimultiple.com/human-in-the-loop-automation>.

¹²² Shobhit Seth and Charles Potters, “The World of High-Frequency Algorithmic Trading,” Investopedia, December 31, 2021, <https://www.investopedia.com/articles/investing/091615/world-high-frequency-algorithmic-trading.asp>.

and implementing HITL imposes additional compliance and governance costs in money and time. Furthermore, because humans may not be able to comprehend the actions of a superhuman stock trader, HITL may prevent advantageous trades that look disadvantageous when overseen by them directly. Generally, there is limited research on when and how to implement HITL to leverage the benefits of robust oversight and capable automation. Both Amazon Web Services⁷⁹ and human-centered AI firm Faculty.ai⁸⁰ highlight the cost of a mistake and the complexity of the decision as key factors in deciding whether to implement HITL in a particular scenario; the cheaper mistakes are or the simpler the problem is, the less necessary human oversight is. The size and quality of the dataset also matter, and still humans may be better at generalizing from and spotting patterns in small, low-quality datasets.⁸¹

HITL also is not a single process; rather there are many different roles the human can take based on the factors above. There is the typical HITL setup in which the human is an active, real-time part of the decision-making process; “human-on-the-loop,” in which the human oversees the AI but does not take a real-time role in decisions; and “human-over-the-loop,” in which they are charged with signing off on the machine’s decisions.⁸² The actions of the human-in-the-loop can also vary from merely triggering automatic processes, to reviewing and verifying outputs in which the AI has low confidence, to intervening and improving future performance.⁸³ Despite its promise, it is important to keep in mind that HITL is not a cure-all for AI mistakes. AI overseers are empirically unreliable in technical oversight due to insufficient competence, harmful incentives,⁸⁴ and even fatigue or boredom.⁸⁵

- B. HOTL: Trading algorithms also are not exempt from making costly errors. Knight Capital Group, then a FINRA-registered broker-dealer, was once one of the largest traders of U.S.

equities; throughout 2011 and 2012, Knight’s aggregate trading represented approximately ten percent of all trading in listed U.S. equity securities.¹²³ However, in August 2012, they deployed a faulty trading algorithm, which, although not AI-powered, had consequences that demonstrate more broadly the necessity of human oversight in automated decision-making. Upon deployment, the algorithm mistakenly placed approximately \$7 billion in orders across more than 150 stocks in less than an hour. Knight requested that their erroneous trades be cancelled, but their request was rejected, and as a result, Knight ultimately realized a net loss of \$460 million.¹²⁴ They never recovered and agreed to be acquired by a rival in December 2012.¹²⁵ If they had better written policies and procedures, better governance and internal supervision, including implementing HITL to reject the anomalous trades, Knight might still exist. Knight Capital Group’s algorithm was an example of high-frequency trading (HFT), a financial trading strategy in which implementing HITL is particularly difficult—due to the high speed and quantity of trades involved—but also particularly important, as the consequences that Knight experienced illustrate.

Without human supervision, mistakes made by HFT algorithms have the potential to create devastating financial losses in a short amount of time. Ideally, HITL should be a best practice for all HFT systems. One possible implementation would be structuring HFT systems to notify a supervisor when the system hits a certain threshold, requiring a manual override to cross the threshold.

¹²³ See, e.g., In re: Knight Capital Americas LLC, Order Instituting Administrative and Cease-and-Desist Proceedings, SEC Release No. 70694 (October 16, 2013), <https://www.sec.gov/files/litigation/admin/2013/34-70694.pdf>.

¹²⁴ Ibid.; John D’Antona Jr., “The Rise and Fall of Knight Capital — Buy High, Sell Low. Rinse and Repeat,” Traders Magazine, September 10, 2019, <https://www.tradersmagazine.com/departments/brokerage/the-rise-and-fall-of-knight-capital-buy-high-sell-low-rinse-and-repeat-2/>.

¹²⁵ Press Release, Knight Capital Group and Getco Holding Company Agree to Merge (December 19, 2012), <https://www.sec.gov/Archives/edgar/data/1060749/000119312512508185/d455870d425.htm>.

Online data privacy

Financial institutions should also take care to respect the privacy of customers' financial data, and behaviors for that matter, particularly in the collection and surveillance of financial information. They should be encouraged to follow proper procedures and compliance with disclosures to the federal government, especially in the face of concerns about national security and financial risk management. To date, the lines are fine between critical surveillance and its impact on consumer privacy. From threats to national security like 9/11 and its aftermath to more invasive financial intrusions by foreign operatives, consumer and institutional data protections can be compromised without appropriate mechanisms to flag concerns.¹²⁶ Voluntary provision of customer financial information is currently legal through the Right to Financial Privacy Act, Section 314(a) of the Patriot Act, and the Bank Secrecy Act.¹²⁷ However, AI software can blur the lines since many technology companies selling their software to the government and other industry partners may have created other 'backdoors' to the direct reporting of compliance information, including demographics, location, etc. Here, the point is that AI companies should be encouraged to come up with clear transparency around the operations and performance of their systems, along with the probable cases and causes for risks.

The talent pipeline

The responsible and trustworthy use of AI tools will also require the creation of a talent pipeline of professionals who are trained in the development and use of AI products. President Biden's Executive

¹²⁶ Sanya Mansoor, "'Who Else Is Spying on Me?' Muslim Americans Bring the Fight Against Surveillance to the Supreme Court," Time, September 16, 2021, <https://time.com/6097712/muslim-american-surveillance-supreme-court-sept-11/>.

¹²⁷ Committee on the Judiciary and the Select Subcommittee on the Weaponization of the Federal Government, "Financial Surveillance in the United States: How Federal Law Enforcement Commandeered Financial Institutions to Spy on Americans," March 6th, 2024, p. 8.

Order on AI and the OMB's memo on AI governance implementation include a long list of requirements for agencies that will need a dedicated workforce to address. As Daniel Ho, Associate Director of the Stanford Institute for Human-Centered Artificial Intelligence and a member of President Biden's National Artificial Intelligence Advisory Committee, told one journalist, current technology leaders within agencies may not have the time or expertise to implement these various AI-related requirements.¹²⁸ Actively recruiting more tech talent and investing in training for the current federal workforce will be crucial in order to meet the goals already set by the Executive Order and the OMB.

AI Applications Among CFTC-Registered Entities

As an agency tasked with promoting the "integrity, resilience, and vibrancy of the U.S. derivatives market through sound regulation," the CFTC has been focused on advances in AI regulation in the financial services sector. And the agency has started to raise awareness among its stakeholders as to the benefits and risks of emerging and evolving technologies, including AI. In September 2023, the CFTC hosted an event on scams and fraud in the digital world and highlighted the recently announced plans of the Federal Trade Commission (FTC) to address AI-powered voice fraud, which can enhance fraudster's ability to deceive their targets at scale.¹²⁹ Overall, however, derivatives trading has not been a focal point of AI regulatory policy thus far, creating room for the CFTC to investigate and address how this increasingly widespread technology will affect the U.S. derivatives market, as well as the destabilizing potential of generative AI-powered tools, such as chatbots, that can super-manipulate markets without guardrails.

¹²⁸ Nihal Krishnan, "AI executive order will face challenges with talent pipeline, experts say," FedScoop, December 7, 2023, <https://fedscoop.com/ai-executive-order-will-face-challenges-with-talent-pipeline-experts-say/>.

¹²⁹ PYMNTS, "CFTC Virtual Event to Focus on Investment Scams and AI," September 25, 2023, <https://www.pymnts.com/fraud-prevention/2023/cftc-virtual-event-to-focus-on-investment-scams-and-ai/>; Security, "FTC announces plans to mitigate AI voice fraud," November 24, 2023, <https://www.securitymagazine.com/articles/100168-ftc-announces-plans-to-mitigate-ai-voice-fraud>.

On this later point, Professor Michael Wellman, who is also a member of the Subcommittee, testified in 2023 to the Senate Banking and Finance Committee about the potential risk of advanced AI in financial markets. In his written testimony, he noted the existing widespread use of AI in algorithmic trading systems.¹³⁰ The latest advances, including generative AI and deep reinforcement learning, extend the scope and autonomy of AI systems and therefore expose new potential vulnerabilities. For example, the impressive fluency of chatbots opens the language channel for manipulation and deception of consumers. This is one example of how AI might produce “super manipulation” capabilities.¹³¹ In the January 2024 TAC Open Meeting, Wellman shared a visual demonstration of his own research where a trading strategy trained by deep reinforcement learning learns to manipulate a financial market benchmark, even though not directed to manipulate by its designer. Machine learning methods could also be used by regulators to detect manipulation, but the same techniques can be used by adversaries to better evade such detection.¹³²

Opportunities of AI Use Among CFTC-Registered Entities

In theory, AI represents a potentially valuable tool for CFTC internal and external stakeholders. This value resides primarily in the potential of AI to improve automated processes governing core functions like risk management, surveillance, fraud detection, and the identification, execution, and back-testing of trading strategies. AI can also be used to collect and analyze information about current or prospective customers and counterparties, and for surveillance and fraud detection. Despite this

¹³⁰ Michael P. Wellman, Written Testimony before the US Senate Committee on Banking, Housing, and Urban Affairs, Hearing on “Artificial Intelligence in Financial Services,” September 20, 2023, https://www.banking.senate.gov/imo/media/doc/wellman_testimony_9-20-23.pdf.

¹³¹ Ibid.

¹³² “Commissioner Goldsmith Romero Announces January 8, 2024, Technology Advisory Committee Meeting,” January 8, 2024, <https://www.cftc.gov/PressRoom/Events/opaeventtac010824>. See also video webcast of meeting with presentations, <https://youtube.com/live/RAwq1m0PVFE>.

potential value, it is difficult to determine whether and how CFTC-regulated firms are currently, or might in the future, use AI. Part of the challenge is a lack of direct knowledge about the organizations currently leveraging AI, in addition to the level of transparency among these firms for regulators and customers, about the components of the AI model, especially as they relate to trading strategies and risk management. But even where firms disclose their use of AI, it is not always clear what type of AI they are using (e.g., generative, or predictive), and whether it is easily distinguishable from other types of algorithmic decision-making or machine learning.

Risks and Challenges

Because of the diverse range of CFTC-registered entities, there could be some difficulties in the identification, measurement, and mitigation of potential risks stemming from the use of AI. Where this AI resembles more conventional forms of algorithmic decision-making, the risks are likely to be remarkably similar and include a heightened risk of market instability, and, especially when combined with high-frequency trading, potential institutional and wider market instability. In this latter respect, the risks stemming from the use of AI by CFTC-registered entities may closely resemble those revealed by the widespread market disruption in August 2012 caused by a software malfunction at market maker Knight Capital.¹³³

The additional risks of generative AI

Meanwhile, where firms use their versions of generative AI, there could be increased risks of institutional and market instability if the CFTC-registered entities themselves do not have a complete understanding of, or control over, the design and execution of their trading strategies and/or risk management programs. The 2023 Annual Report of the Financial Stability Oversight Council pointed to

¹³³ See *supra* note 123–125.

this cognitive gap between the generative AI models and the CFTC registrants adopting those models: “With some generative AI models, users may not know the sources used to produce output or how such sources were weighted, and a financial institution may not have full understanding or control over the data set being used, meaning employment of proper data governance may not be possible.”¹³⁴

If “proper data governance” is not currently possible for some registered entities employing generative AI models, the CFTC may have to propose some rules or guidance to enable firms to access the AI models in a semi-autonomous way that will balance the intellectual property protection of the AI model provider with the imperative of the firm using the model to properly manage and report its trading and clearing data. That is why having a functional taxonomy of the objectives of CFTC-registered entities may be a useful starting point to future actions, including enforcement. Understanding the AI use cases of CFTC-registered entities will enable regulators to define proper and improper use for these and other scenarios.

V. AI Risks and Management Among CFTC-Registered Entities

AI Risks

Certain AI risks are at this point widely identified and well-known among those working with AI systems in industry, government, and academia.¹³⁵ These risks include those related to the lack of transparency or explainability of AI models’ decision-making processes (i.e., the “black box” problem).

¹³⁴ Financial Security Oversight Council, “Annual Report 2023,” 2023, <https://home.treasury.gov/system/files/261/FSOC2023AnnualReport.pdf>, p. 94.

¹³⁵ E.g., NIST AI RMF 1.0 at 4-9; “Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence,” Office of Management and Budget, November 1, 2023, <https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf> pp. 24-25; “AI risk atlas,” IBM, March 21, 2024, <https://www.ibm.com/docs/en/watsonx-as-a-service?topic=ai-risk-atlas>; Bommasani et al., “On the Opportunities and Risks of Foundation Models,” arXiv, July 12, 2022, pp. 130-59, <https://arxiv.org/abs/2108.07258>.

There are also issues with data relied on by AI systems, including overfitting of AI models to their training data or the “poisoning” of real-world data sources encountered by the AI model during its deployment.¹³⁶ Additionally, there is potential mishandling of sensitive data, including personal identifiable information (PII), within AI systems; fairness concerns, including the possibility that an AI system could reproduce or compound biases it encounters in its training data, in its use-specific adaptation data, or in real-world applications. At times, there are concentration risks that arise from the most widely deployed AI foundation models relying on a small number of deep learning architectures, as well as the relatively small number of firms developing and deploying AI foundation models at scale. Lastly, there are risks related to AI systems' potential to produce false or invalid outputs, whether because of AI reliance on inaccurate “synthetic” data to fill data gaps or because of unknown reasons (e.g., hallucinations). Numerous other risks have also been identified in the emerging field of AI risk management.

More specific areas of risk can be identified by considering these general areas of AI within the context of particular use cases likely to occur in CFTC-regulated markets. It should therefore be a central focus for the CFTC and relevant risk management professionals in the industry to identify the more specific risks that have high salience for CFTC-regulated markets, and then measure the potential harm that could occur should these risks be insufficiently managed. To aid these efforts, this Subcommittee identified a partial list of use cases among CFTC-registered entities, and areas of risk likely to be relevant. A more expanded table of specific use cases is included in Appendix One.

¹³⁶ See generally “Managing model and AI risks in the investment management sector,” Deloitte, 2023, <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/Advisory/us-advisory-deloitte-managing-model-and-ai-risks-in-the-investment-management-sector-july-2023.pdf>; Zhang, *supra* n. 116.

- **Trading and investment use cases:** Investment research and analysis including the analysis of large amounts of unstructured data; identification of optimal trade execution strategies (e.g., timing and venue); predicting asset prices by uncovering complex and hitherto unidentified correlations; algorithmic and high-frequency trading (e.g., improving parameters and decision logic of existing trading algorithms, developing and modifying trading algorithms based on natural language instructions, creating “self-correcting” algorithms that improve over time, etc.).
 - **Relevant areas of risk:** Data overfitting (e.g., model overfitting to historical trade or pricing data); data poisoning; business continuity risks posed by dependence on a small number of AI firms (i.e., critical infrastructure dependence risk); procyclicality risks or other risks caused by multiple firms deploying similar AI models in the same market; erroneous AI output or decisional errors causing sudden, substantial losses to a particular firm, asset class, or market (i.e., a “flash crash”).

- **Customer communications, advice, and service use cases:** Marketing; customer acquisition; customer profiling and retention; customer service chat and call functions; personalized investment or trading advice.
 - **Relevant areas of risk:** AI hallucinations or corrupted data causing inaccurate AI outputs (e.g., false or misleading marketing materials or individual customer communications); privacy risk related to AI systems processing and using the sensitive PII of customers; explainability and transparency risks (e.g., potential inability to provide a rationale demonstrating AI communications met fiduciary duty standards or tailored risk

disclosure requirements); bias and discrimination in treatment of customers, provision of financial advice, or placement of marketing materials.

- **Risk management use cases:** Margin model optimization and monitoring; hedging strategy development and execution; monitoring and adjustment of excess funds requirements; monitoring unstructured data (e.g. news, social media, etc.) for developments related to key depositories, counterparties, third-party service providers, etc.; collateral and liquidity optimization.
 - **Relevant areas of risk:** Data quality and data poisoning (e.g., if an AI system monitoring news sources is unable to distinguish more reliable and less reliable sources, or if false news stories are maliciously published to influence AI systems); AI hallucinations; critical infrastructure dependency risks; bias and discrimination (e.g., if an AI system consistently calculates higher margin levels for minority-owned counterparties); explainability (e.g., if a logical rationale for adjusting or failing to adjust a margin model cannot be given).
- **Regulatory compliance use cases:** Market surveillance; recordkeeping and reporting; KYC compliance in customer or counterparty onboarding; AML/CTF transaction monitoring; regulatory capital and margin calculations and optimizing the allocation of regulatory capital and margin; monitoring employee communications; monitoring transfers and withdrawals from customer and proprietary accounts.
 - **Relevant areas of risk:** Critical infrastructure dependence; AI hallucinations; bias and discrimination; explainability (e.g., lack of audit trail on transaction monitoring decision-making process); data privacy.

- **Back office and operations use cases:** Automation in trade verification, reconciliation, reporting, etc.
 - **Relevant areas of risk:** Critical infrastructure dependence; AI hallucinations; explainability; data privacy.

While broad, and in some instances, overlapping, this initial detail of AI risks should be expanded on by the CFTC through planned roundtables, or convenings with registered entities and other stakeholders.

AI Risk Management Framework for CFTC-Registered Entities

CFTC registrants are required to manage the risks associated with their business. This pre-existing regulatory framework, which is quite broad in some respects, likely already reaches many of the AI-associated risks that have been discussed throughout this report. In other areas, CFTC regulations may need to be clarified through staff guidance or amended through Commission rulemaking to require firms to manage AI risks appropriately. Given the unique nature of AI risks and the challenges firms are likely to face in managing these risks, this section recommends that the CFTC conduct a consultation to help establish AI governance standards before instituting specific requirements as to AI risk management. Whether existing regulations already address AI-related risks, or new or amended regulations are needed, significant questions remain regarding how to implement effective risk management and oversight of AI systems at the firm level, questions that this Subcommittee believes can be best approached in the first instance by establishing good governance over AI. In keeping with the Commission's longstanding practices, the CFTC should ensure that any framework implemented is principles-based without sacrificing rigor, to allow for flexibility as the use of advanced AI becomes more prevalent in CFTC-regulated markets.

Existing Risk Management Requirements

Turning to existing risk management requirements, futures commission merchants (FCMs) and swap dealers (SDs) are both required to have a risk management program (RMP).¹³⁷ While there are important distinctions between the FCM and SD-specific RMP regulations, at a high level both regulations require the relevant firms to have in place written policies and procedures sufficiently designed and tailored to monitor, manage, and supervise the risks of their activities and take into account specific risk management considerations; to maintain an internal risk management unit with qualified personnel, sufficient authority and resources, and independence from the business unit; and to periodically submit risk exposure reports to the relevant regulators.¹³⁸ As a governance best practice, FCMs and SDs should, on their own accord, implement and test tailored written policies and policies addressing AI-related risks to their specific business. However, FCM and SD risk management regulatory requirements arguably already cover many AI-related risks, albeit in different ways.

FCM risk management requirements are generally codified in CFTC Regulation 1.11. Under this regulation, an FCM must consider “technological” risks as part of its RMP, which would presumably cover an FCM's use of AI systems in conducting its business as an FCM (e.g., any use of AI in core functions related to trading, margining, monitoring its capital and margin and segregated funds, etc.).¹³⁹ While this is a very high-level requirement, some of the more specific aspects of the FCM risk management regulation also would appear to affect the use of AI systems in that particular area. For

¹³⁷ Another category of market intermediary, retail foreign exchange dealers (RFEDs), have a small number of active registrants. RFEDs have a “risk assessment” requirement that is more limited than the requirements imposed on FCMs and SDs. See generally 17 C.F.R. § 5.10. However, as of the writing of this report, nearly all RFEDs are dual-registered as FCMs and are therefore bound by the more stringent FCM RMP requirements.

¹³⁸ See generally 17 C.F.R. § 1.11(c), (d); § 23.600(b).

¹³⁹ See 17 C.F.R. § 1.11(e)(1).

example, in managing operational risk, an FCM must ensure it has policies and procedures over “automated trading programs” to govern the use, supervision, maintenance, testing, and inspection of such automated programs.¹⁴⁰ The requirement to supervise and reassess the adequacy of excess funds target levels would also presumably include a review of use of AI systems in setting those levels. However, many AI-specific risks are not addressed by the FCM risk management regulatory requirements.

SD risk management requirements are set out in CFTC Regulation 23.600, and many if not all areas of AI-related risk are addressed within the CFTC's enumerated risk management considerations, including credit and operational risk. For example, an SD's operational risk policies and procedures must consider: the security and reliability of operating and information systems, as well as their scalability and independence from the business unit; safeguards to detect, identify and promptly correct deficiencies in operating and information systems; and the reconciliation of all data and information in operating and information systems.¹⁴¹ These areas of operational risk accord with key AI-related risks described in this report, including risks regarding data quality, explainability, and bias (in this case, the potential for AI systems to become improperly biased in favor of business-unit concerns). SDs also have broad requirements to manage the risks associated with new products and with the SD's business trading unit¹⁴²; this would presumably encompass the obligation to manage any AI-related risks associated with new products and AI-directed or AI-advised trading.

¹⁴⁰ See 17 C.F.R. § 1.11(e)(3)(ii).

¹⁴¹ See 17 C.F.R. § 23.600(c)(4)(vi).

¹⁴² See 17 C.F.R. § 23.600(c)(3); 23.600(d).

For exchanges, clearinghouses, and data repositories (a group collectively referred to as “registered entities”¹⁴³), the CFTC regulates these entities through a core principles regime that provides the registrant with significant discretion in how to meet the CFTC's requirements. Relevant here, registered entities are subject to largely identical regulations in respect to the “System Safeguards” core principle.¹⁴⁴ Under this core principle, registered entities must have in place a program of risk analysis and oversight that includes the development of appropriate controls.¹⁴⁵ “Controls” are in turn defined to include safeguards that protect the reliability, security or capacity of the registered entity's automated systems, or the confidentiality, integrity, or availability of the registered entity's data.¹⁴⁶ As with the definition of “operational risk” discussed above as to SD risk management requirements, this broad language regarding controls appears on its face to capture many AI-related risks such as data quality, data overfitting, hallucinations and other types of erroneous output, and explainability.

Thus, many AI-related risks seem to be addressed by existing CFTC risk management requirements. However, more questions about AI risk management remain and would benefit from more engagement from the CFTC – as well as a gap analysis.

Paths Forward for AI Risk Management in CFTC-Regulated Markets

A central question that the CFTC could consider is whether and to what extent firms should incorporate an AI-specific risk-management function into their existing risk management structures. As described throughout this section, many AI-related risks are likely already captured to the extent they impact the integrity of operational systems and data, cybersecurity, margin and capital requirements,

¹⁴³ 17 C.F.R. § 1.3. The list of registered entities includes designated contract markets (DCMs), swap execution facilities (SEFs), derivatives clearing organizations (DCOs), and swap data repositories (SDRs).

¹⁴⁴ See generally 17 C.F.R. Part 37, Subpart O (SEF); Part 38, Subpart U (DCM); § 39.18 (DCO); § 49.24 (SDR).

¹⁴⁵ 17 C.F.R. § 37.1400(a) (SEF); § 38.1050(a) (DCM); § 39.18(b)(1) (DCO); § 49.24(a)(1) (SEF).

¹⁴⁶ 17 C.F.R. § 37.1401(h)(1) (SEF); § 38.1051(h)(1) (DCM); § 39.18(a) (DCO); § 49.24(j)(1) (SEF).

and so on. The CFTC could perhaps proceed simply by adding AI-specific risks not clearly covered by existing regulations (such as risks posed by embedded biases in AI training data, or risks posed by reliance on AI technology firms outside of the CFTC's regulatory jurisdiction) to the lists of risks that existing risk management programs must consider. In one sense, this seems appropriate—as NIST has noted, AI risks should generally be addressed as part of a broader enterprise risk management program.¹⁴⁷ Yet at the same time, AI technology and especially the most advanced AI models represent a novel and unique source of risk that may warrant more intensive risk-management efforts from those firms deploying AI technology.

Additionally, simply adding certain AI-specific risks to current risk-management requirements may be more difficult than it would at first seem. Adequate risk management of AI systems will likely need to involve teams with broad and deep expertise from different departments across a firm (e.g. risk management, compliance, business, and information systems), including individuals who participate in back, middle, and front-office activities. But while this approach is likely necessary for adequate risk management,¹⁴⁸ it may also pose challenges for maintaining the independence of a firm's risk-management function in respect to AI, as the interests of various departments including the business unit could seek to influence or override risk management decision-making. A related issue likely to arise in AI-specific risk management is how and where in the AI system lifecycle to require human oversight or even direct human intervention in AI processes, as well as who within the firm should perform these functions and what their qualifications should be—i.e., the HITL question. A third set of key questions for AI risk management requirements to address will revolve around vendor risk, including issues related

¹⁴⁷ NIST, “Artificial Intelligence Risk Management Framework 1.0,” p. 8, January 2023, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.

¹⁴⁸ See, e.g., *ibid* p. 26 (citing the need for interdisciplinary teams).

to information access, potential lack of control, and transparency surrounding data inputs and data retention; widely varying capacity levels at CFTC-registered entities; and unequal bargaining power between CFTC-registered entities and the large technology companies offering AI models for business use. Specific answers to all these risk management questions are likely to vary by firm and by the particulars of how different AI models are deployed in different business contexts.

Therefore, considering the specificity and variability of AI use cases at CFTC-registered entities, and the potentially transformative nature of advanced AI models, it likely makes more sense for the CFTC to avoid simply imposing additional enumerated AI-related risks to existing risk management requirements (at least initially) and instead to develop appropriate firm-level *governance standards* over AI systems. This would accord with the NIST Framework, which recommends that AI risk management begin by establishing appropriate governance.

At a high level, and following the NIST Framework, AI governance would likely involve first putting in place policies and procedures related to mapping, measuring, and managing AI risks, and should be keyed to the core features of responsible or ethical AI. Appropriate AI governance also means ensuring that individuals with responsibility for AI systems oversight have appropriate qualifications and training. Decision-making relating to AI should be informed by a diverse team (which NIST notes means not only diverse in terms of demographics but also in terms of expertise, disciplines, and so on). Appropriate governance standards must also provide for robust engagement with relevant AI actors both inside and outside of the firm: firm personnel who have the closest interaction with AI systems must have clear reporting responsibilities supported by a safety-first culture, and their feedback should be collected, considered, and prioritized. There should also be communication and feedback processes for outside personnel who have actionable knowledge of the development, foundational training, and adaptive training of the relevant AI systems. Finally, governance standards for AI should address risks

arising from the involvement of and/or concentrated reliance on outside service providers in a CFTC registrant's core business or regulatory functions (e.g., risks related to any potential disruption of a critical AI technology firm's business operations, whether due to cybersecurity events or other causes; and transparency risks associated with such a critical firm's likely status of being outside the CFTC's jurisdiction).

VI. Recommendations

This Report has surfaced various governance, design, and accountability concerns that the CFTC might consider as the use of AI evolves and increases among its relevant stakeholders. Obviously, there is an urgent and timely need for the Commission to engage further into some of the areas discussed, including the future applications of generative AI, and other frontier models on its regulated firms and exchanges. And given the complexity of the topic, the agency is encouraged to involve more stakeholders to level set definitions and implications of AI technologies. At the core of any CFTC engagement should be a framework that balances risk, responsible innovation, and the need for more ethical practices and applications. The latter is currently being debated, and as mentioned, some regulatory agencies like those in the EU have drafted clear guidance and, in some instances, more prescriptive regulation on how to achieve this.

What the Subcommittee realized as the process evolved is that there are various definitions and understandings of AI based on where one sits in the broader ecosystem. Many of the CFTC-registered entities may be unaware about how existing and emerging public policies could impact their business models and practices.¹⁴⁹ Thus, the Subcommittee encourages the CFTC to start with baseline

¹⁴⁹ U.S. Department of the Treasury, "Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Sector," March 2024, <https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf>.

engagement for the organizations under their jurisdiction to ensure that they understand the future implications of AI, and evolving technologies like generative AI.

In the concluding section, the Subcommittee offers a series of proposed recommendations with some explanation as to whom the specific action is directed, and the extent to which the proposal is one of awareness or more prescriptive posturing. In some of the recommendations, we urge the CFTC to leverage its role as a market regulator to support the current efforts on AI coming from the White House and Congress.

During our meetings, the Subcommittee also realized that a more in-depth conversation must be had around the current and future use cases of generative AI, as well as any potential erosion in confidence or aversion to risk that this technology holds. It is also imperative that the CFTC focus in on the opportunities and drawbacks of LLMs, which are being deployed for very basic administrative capacities in the financial services industry including customer chatbots, to those that are driving trading and other more significant institutional capabilities. Here is yet another area where the CFTC can offer granular and global guidance on generative AI tools, and the interpretation of outcomes.

The remainder of the section shares the proposed recommendations of the Subcommittee to address some of the concerns amplified in the report: (1) to support the Commission's stance on the future use of AI for a variety of functions cutting across administration to organizational client management; (2) to raise awareness about the more appropriate use cases for CFTC-related workstreams to avert overfitting of AI and other highly capable models; and (3) to align the Commission with the adjacent and related activities around AI governance that will ultimately impact the agency's current and future work flows. While there is still much more to be discovered and explored with the proliferation of AI among entities regulated by the CFTC, the Subcommittee suggests that the

Commission and the specific role of the TAC start to develop a framework that fosters safe, trustworthy, and responsible AI systems, and consider the five *Proposed Recommendations* below.

Proposed Recommendations

A. Recommendation One: The CFTC should host a public roundtable discussion and CFTC staff should directly engage in outreach with CFTC-registered entities to seek guidance and gain additional insights into the business functions and types of AI technologies most prevalent within the sector.

- *Intended purpose*: To inform the CFTC about key technical and policy considerations for AI in financial markets, develop common understanding and frameworks, build upon the information of this report, and establish relationships for future discussion. The discussion topics should include, but not be limited to, humans-in-or-around-the-loop of the technology, acceptable training data use cases, and development of best practices and standards as it relates to the role of AI.
- *Intended audiences*: CFTC regulators, staff, registered entities, and third parties who can use insight to design appropriate AI regulatory policy and/or best practices.
- *Potential outcomes/deliverables*: Roundtable discussion and supervisory discussions and consultations with CFTC-registered entities to ascertain how AI systems are used in markets and how future AI developments may impact markets.

B. Recommendation Two: The CFTC should consider the definition and adoption of an AI Risk Management Framework (RMF) for the sector, in accordance with the guidelines and governance aspects of the NIST, to assess the efficiency of AI models and potential consumer harms as they apply to CFTC-registered entities, including but not limited to governance issues.

- *Intended purpose*: To ensure some certainty, understanding and integration of some of the norms and standards being developed by NIST, and to introduce these practices to regulated industries and firms.
- *Intended audiences*: CFTC regulators, staff, registered entities, and others who will be impacted by the AI RMF, and its intended outcomes.
- *Potential outcomes/deliverables*: A potential proposed rule from the CFTC applying the information from Recommendation One to implement the NIST framework, thus ensuring financial markets and a regulatory system that is more resilient to emerging AI technologies and associated risks.

C. Recommendation Three: The CFTC should create an inventory of existing regulations related to AI in the sector and use it to develop a gap analysis of the potential risks associated with AI systems to determine compliance relative to further opportunities for dialogue, and potential clarifying staff guidance or potential rulemaking.

- *Intended purpose*: To confirm CFTC's oversight and jurisdiction over increasingly autonomous models, and to make more explicit compliance levers.

- *Intended audiences*: CFTC regulators, staff, registered entities, and others who have interest in AI's compliance.
- *Potential outcomes/deliverables*: Issue clarifying staff guidance or proposed potential rulemaking from the Commission that advances the explicit and implicit applications of existing regulatory measures.

D. Recommendation Four: The CFTC should strive to gather and establish a process to gain alignment of their AI policies and practices with other federal agencies, including the SEC, Treasury, and other agencies interested in the financial stability of markets.

- *Intended purpose*: To leverage and utilize best practices across agencies, and potentially drive more interagency cooperation and enforcement.
- *Intended audiences*: CFTC regulators, staff, and other similarly aligned regulatory agencies.
- *Potential outcomes/deliverables*: Interagency meetings and cooperation.

E. Recommendation Five: The CFTC should work toward engaging staff as both 'observers' and potential participants in ongoing domestic and international dialogues around AI, and where possible, establish budget supplements to build the internal capacity of agency professionals around necessary technical expertise to support the agency's endeavors in emerging and evolving technologies.

- *Intended purpose*: To build the pipeline for AI experts at the agency, and to ensure necessary resources for staff, events, and other related activities

that ensure more responsible engagement of AI by internal and external stakeholders.

- Intended audiences: CFTC regulators and staff.
- Potential outcomes/deliverables: Increased budgetary resources toward AI services, and more presence at relevant conferences and convenings.

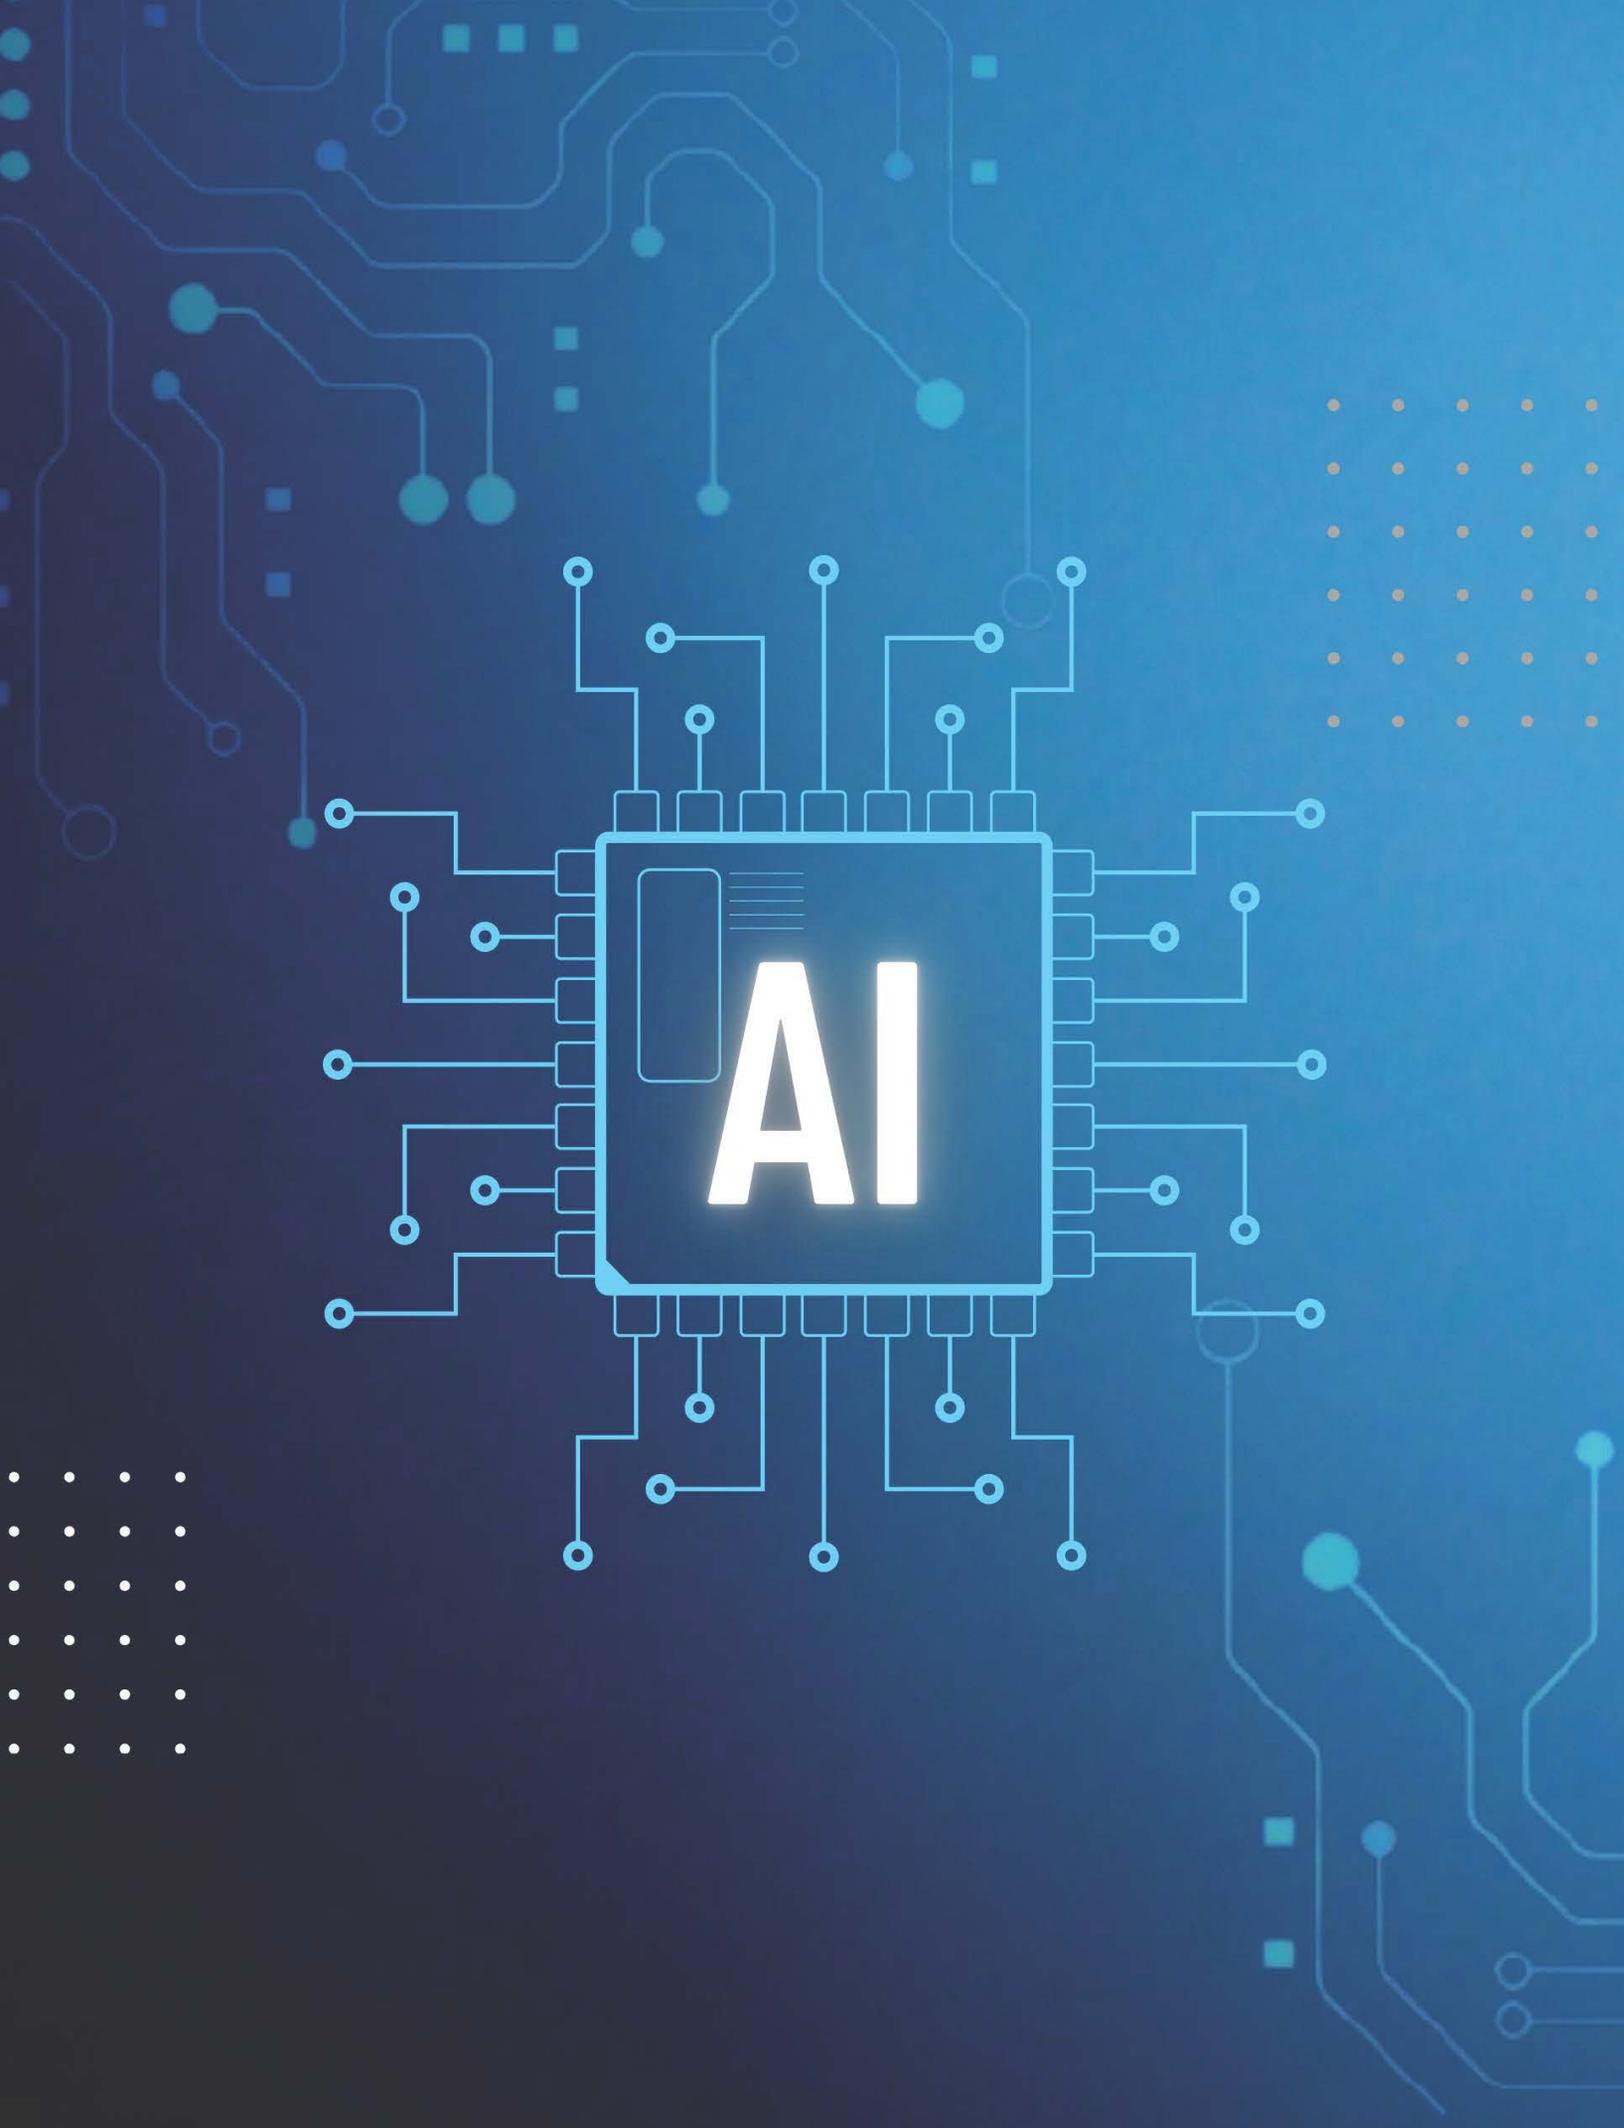
These recommendations are also presented with the intended audiences for each, and potential deliverables. In some of the recommendations, we urge the CFTC to leverage its role as a market regulator to support the current efforts on AI coming from the White House and Congress.

APPENDIX ONE: ANNOTATED RISK TABLE

Risk taxonomy of defined AI use cases and functions for CFTC-registered entities

Functions	Use Cases	Potential Risks
Trading and investment	<ul style="list-style-type: none"> • <i>Investment research and analysis.</i> AI can support greater automation in the collection, cleaning, manipulation, and analysis of data. NLP tools can be used to process large volumes of diverse, unstructured, and alternative data to reveal hidden patterns, trends, anomalies, and insights. • <i>Trading execution strategies.</i> AI can be used to develop trading strategies and identify optimal order execution: e.g. timing and venues. • <i>Asset pricing.</i> AI can be used to predict asset prices using historical and cross-sectional data, uncovering non-linear interactions among numerous variables. • <i>Asset allocation.</i> AI can be used to identify the optimal allocation of resources among different asset classes to maximize return on investment. This includes the development of dynamic asset allocation strategies. • <i>Algorithmic trading:</i> AI can be used to improve the parameters and decision logic used in connection with algorithmic trading models. • <i>High-frequency trading.</i> AI, including transformer and deep learning models, can be combined with algorithmic trading strategies to identify trading opportunities and execute trades extremely rapidly. 	<ul style="list-style-type: none"> • <i>Data quality.</i> AI models may be unable to evaluate the credibility of data and data sources. • <i>Data overfitting.</i> AI models can become too specialized in the training data and fail to generalize well to new data. • <i>Explainability.</i> AI models are highly complex and may not be capable of comprehension by human users. This makes it difficult or impossible for users to fully understand or describe the relevant inputs, decision-making process, or outputs. • <i>Firm-level knowledge and expertise.</i> Regulated firms will need to recruit and retain professionals with sufficient expertise in AI in order to ensure effective design, training, and oversight. • <i>Critical infrastructure dependence.</i> As regulated firms become more reliant on a small number of AI vendors and platforms, disruptions in service may have knock-on effects for these firms and their ability to carry on business, manage risks, and comply with regulatory requirements. • <i>Procyclicality.</i> The widespread use of a small number of similar AI tools and/or data sources could contribute to correlated trading strategies and the convergence of industry practices. Highly correlated trading strategies and industry practices could trigger or amplify market disruptions, contributing to procyclicality and potential financial instability. • <i>Algorithmic trading and market disruption.</i> Where AI is combined with algorithmic trade execution, the failure at the appropriate level (e.g., exchanges, intermediaries, etc.) to implement appropriate circuit breakers can leave regulated firms vulnerable to losses in disorderly markets. The automated execution of these trades can also contribute to disorderly markets, generating self-reinforcing feedback loops.
Customer advice and service	<ul style="list-style-type: none"> • <i>Customer acquisition and marketing.</i> AI can be used to enhance marketing, lead generation, and customer segmentation. 	<ul style="list-style-type: none"> • <i>Data quality.</i> AI models may be unable to evaluate the credibility of data and data sources. AI models can also fill data gaps with synthetic data, leading to inaccurate customer information.

	<ul style="list-style-type: none"> • <i>Customer profiling.</i> AI can be used to collect and analyze large volumes of data from multiple sources to generate insights about customer profiles, activity, and preferences. • <i>Customer service.</i> AI can be used to automate customer-facing interactions, either supporting or replacing human relationship managers. • <i>Financial advice.</i> AI can be used to analyze customer data to generate personalized product recommendations, including potential investment opportunities. 	<ul style="list-style-type: none"> • <i>Customer privacy.</i> AI models may collect data from sources without regard to compliance with relevant customer privacy laws. • <i>Implicit bias and discrimination.</i> AI models trained on data that reflects implicit biases may discriminate against certain classes or categories of customers. • <i>Fiduciary duties (explainability).</i> Where regulated firms owe fiduciary duties to their customers, the lack of explainability associated with AI models may interfere with the discharge of these duties. • <i>Critical infrastructure dependence.</i> As regulated firms become more reliant on a small number of AI vendors and platforms, disruptions in service may have knock-on effects for these firms and their ability to carry on business, manage risks, and comply with regulatory requirements.
Risk management	<ul style="list-style-type: none"> • <i>Hedging.</i> AI can be used to develop, execute, and continuously update hedging strategies designed to maintain a desired risk-return profile. • <i>Collateral and liquidity optimization.</i> AI can be used in collateral and liquidity optimization to forecast liquidity requirements, assist in determining when and how to move collateral in support of trading and settlement, or offer guidance on the selection of eligible collateral. 	<ul style="list-style-type: none"> • <i>Data quality.</i> AI models may be unable to evaluate the credibility of data and data sources, which could lead to poisoning of the models and result in automated taking of suboptimal or detrimental actions in response to real or perceived risks. AI models can also fill data gaps with synthetic data, leading to inaccurate information. • <i>Critical infrastructure dependence.</i> As regulated firms become more reliant on a small number of AI vendors and platforms, disruptions in service may have knock-on effects for these firms and their ability to carry on business, manage risks, and comply with regulatory requirements.
Regulatory compliance	<ul style="list-style-type: none"> • <i>Market surveillance.</i> AI can be used to detect market manipulation, fraud, and unfair practices. • <i>Regulatory capital and margin.</i> AI can be used in connection with the calculation and optimization of regulatory capital and margin and other regulatory requirements. • <i>KYC/AML/CTF.</i> AI can be used in connection with customer onboarding (KYC) and AML-CTF detection. 	<ul style="list-style-type: none"> • <i>Critical infrastructure dependence.</i> As regulated firms become more reliant on a small number of AI vendors and platforms, disruptions in service may have knock-on effects for these firms and their ability to carry on business, manage risks, and comply with regulatory requirements.
Back office and operations	<ul style="list-style-type: none"> • <i>Back office processes.</i> AI can support greater automation and efficiency in back office processes including trade verification, reconciliation, and reporting. 	<ul style="list-style-type: none"> • <i>Critical infrastructure dependence.</i> As regulated firms become more reliant on a small number of AI vendors and platforms, disruptions in service may have knock-on effects for these firms and their ability to carry on business, manage risks, and comply with regulatory requirements.

A central graphic of a square chip with the letters 'AI' in white. The chip is surrounded by a network of white lines representing circuitry, with various nodes and connections. The background is a gradient of blue, with faint circuit patterns and a grid of orange dots on the right side.

AI