coinbase

To:
Moses Kim
Director, Office of Financial Institutions Policy
U.S. Department of the Treasury
1500 Pennsylvania Avenue, NW
Washington, DC 20220

August 12, 2024

Re: Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Coinbase, Inc. (**CBI** and, together with its subsidiaries, **Coinbase**) appreciates the opportunity to respond to the Request for Information (the **RFI**) published by the Department of Treasury's Office of Financial Institutions Policy (**Treasury**).

Coinbase started in 2012 with the idea that anyone, anywhere, should be able to send and receive Bitcoin easily and securely. Today, we are publicly listed in the United States and provide a trusted and easy-to-use platform that millions of verified users in over 100 countries rely on to access the crypto economy.

There is a significant opportunity to improve the financial services sector by leaning in to the use of AI. The digitization of financial products and the resulting automation of their trading has resulted in the creation of staggering amounts of data on a daily basis. This data can be used to train effective AI models that can then in turn analyze market data to enhance market integrity and improve the consumer experience. Blockchain technology can make these systems even better.

Nonetheless, the adoption of AI by financial institutions must be done responsibly. Treasury and other financial services regulators should partner with financial institutions to make sure that appropriate protections are in place while not stifling innovation.

Coinbase appreciates Treasury's engagement with this issue which we expect to only become more important in financial services. We look forward to continuing to work with you as this technology develops.

Yours sincerely,

Faryar Shirzad
Chief Policy Officer, Coinbase

Introduction

Coinbase applauds the release of the RFI because it demonstrates concretely that Treasury appreciates the impact that machine learning technology ("ML")[1] will have—and is already having—on financial markets and the financial services industry. Indeed, Coinbase believes that ML methods will be an important tool for growth, innovation, and consumer protection in the financial services sector. Today, most of our experience with the ML methods underlying generative AI products and processes is associated with supervision of market activities in CBI.

Greater adoption of AI-enabled technologies like those that we and other market participants are adopting will not only enable regulated entities to better meet their own regulatory obligations but will also help financial services regulators fulfill their missions. In particular, we believe that regulators and self-regulatory organizations that use AI-based systems responsibly for purposes of fraud prevention and deterrence of market manipulation and other potential abuses will more effectively ensure orderly markets and investor protection than those that do not. As a consequence, these entities will enjoy higher integrity and offer greater safety for investors and consumers.

To that end, regulators should avoid adopting expansive rules that risk stifling innovation and investment in this nascent technology.[2] Laws that create complex or potentially unclear rules for the use of AI—like those we are already seeing in other jurisdictions—may inhibit innovation by favoring a small set of large firms with the resources necessary to successfully navigate them. This would serve to entrench incumbent entities while limiting contributions from more innovative newcomers. The promise of AI in this area cannot be realized in an overly-prescriptive regulatory environment that limits the ability of entities to embrace future developments.

To be sure, the use of AI systems and tools is already present in the financial services sector. ML, which is captured in the definition of AI both by the RFI and by President Biden's Executive Order No. 14110,[3] is used by Coinbase in certain carefully governed ways to improve processes within the company, with promising results. Beyond ML, systems using AI, as defined by the Executive Order, are not yet widely deployed by

[1] Machine learning defines a related but distinct area of computer science that focuses on the development and use of algorithms that enable computers to learn from and make predictions or decisions based on data without being explicitly programmed for each task.

[2] *See, e.g.*, *EU AI Act: First Regulation on Artificial Intelligence, European Parliament* (June 8, 2023), https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence (last updated June 18, 2023).

[3] *See* Exec. Order No. 14110 (Oct. 30, 2023), 88 Fed. Reg. 75191, at 75193, 75195 (Nov. 1, 2023) (defining "machine learning") https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf; RFI at 5 (referring to "innovations in AI, including machine learning.").

Coinbase for product development and production, risk management, or other corporate functions.

Market Surveillance and Integrity

The most prevalent use of ML systems at Coinbase is for trade surveillance programs on CBI and Coinbase Derivatives, LLC (CDE) platforms. As we have expressed in other consultation responses, these systems can be designed to detect manipulative trading activities, such as “spoofing” or “layering,” by observing trade message patterns that indicate such schemes.[4] These ML surveillance models can learn over time (through programmatic evolution as well as input from surveillance team analysts) which data patterns should trigger a regulatory alert to the market operator consistent with its surveillance and investigatory policies and procedures.

CBI is already using ML models to assist trade surveillance to reduce the escalation of false positives. The ML models deployed for CBI assign a probability score that is generated using fixed inputs. The Surveillance team sets the automation logic to close all alerts below a probability threshold score and to escalate for human review those above the Surveillance defined score. Procedure parameter settings for manipulative activities such as spoofing and layering are initially set to be conservative so that alert scores result in a high number of false positives and regulatory alerts being generated. Guided by the Surveillance staff’s probability score threshold setting, the ML model will auto-close a majority of the false positives and allow analysts to focus on more high-probability manipulative activity.

These same methods are also now on track to be used by our derivatives exchange CDE as part of our self regulatory obligations, to assist in fine tuning alert parameters and analyzing market data and participant activity. ML stands to provide a new perspective into how CDE’s markets and participants operate, as well as potentially identifying and highlighting new disruptive practices.

CDE and CBI have observed substantial efficiency gains in running their respective trade surveillance programs with these tools. These techniques enable 24/7/365 monitoring across all of Coinbase’s trading platforms in a way that manual tracking alone cannot match. Unlike traditional market surveillance, which is often done forensically after the fact, these tools help to provide our Trade Surveillance teams with real-time insights that can be actioned and, often, mitigated quickly.

[4] Coinbase, *Re: Request for Comment on the Use of Artificial Intelligence in CFTC-Regulated Markets*, (April 24, 2024), *available at* https://assets.ctfassets.net/c5bd0wqjc7v0/4Hodz91tFI3ggSTBBBLR8r/641c68aeff2d28c46e618206506d3a3d/CFTC_Response.pdf.

Particularly, as an ML model is trained to reduce false-positive alerts, the number of regulatory alerts processed and requiring review can be reduced over time. Similarly, the number of alerts that require further escalation and review also can be reduced. Designed and programmed responsibly, with the appropriate level of human intervention and other redundancies (including a robust quality control review program and model validation procedures), ML models can assist with the scaling of trade surveillance programs while at the same time improving their efficacy.

Improving the Customer Experience

Coinbase entities are also either using ML models or exploring the use of AI systems to improve the overall customer experience. AI and ML models hold great promise in improving the customer onboarding process, including in particular the important step of verifying customer identification. In addition, as discussed above, these tools are well-suited for detecting and preventing fraud and manipulation on our platform.

Customer Onboarding

Verifying customer identity is a fundamental precursor to setting up a CBI account. CBI's customer identification program has policies and procedures that enable the company to confirm the identity of potential customers seeking to onboard. These same policies govern the use of third-party vendors that assist with customer identification.

The onboarding process involves asking for personal information, as required by applicable regulations. That work generates a risk score for the customer. Based on that score, the customer may also be required to undergo "Enhanced Due Diligence" ("EDD"), where Coinbase may request additional information, such as information about the customer's source of funds, to determine if the customer should have access to the Coinbase platform. The customer risk score is also dynamic. For example, a customer that was not subject to EDD during onboarding may be subject to EDD at a later time based on their platform activity.

The ID-verification process also involves the use of software that confirms the veracity of submitted documentation and its association with the onboarding customer through a variety of different methods. It is often during this onboarding stage where instances of first-party fraud attempts arise, which is where a person knowingly attempts to misrepresent their identity or give false information for financial or material gain.

Increasingly, ML systems can assist with the ID-verification process and mitigate first-party fraud. CBI has leveraged ML models to further automate the onboarding process and reduce the risks of human error that might enable fraudulent behavior. For example, an ID-verification model can be designed where an ML program can ingest photo images of the onboarding customer provided by the customer and then process those images by comparing them to a real-time photo taken of the customer during the onboarding process, as well as other facial images found elsewhere in the public domain. Additionally, such a model could ingest and process other verifying documentation provided in order to tie the customer's personal information to other data available and provided to the model.

Collectively, all of this data concerning one onboarding provides a significant amount of information that can be leveraged to create a safer, more compliant experience. The model also can be programmed to identify or flag any anomalous data for additional review, or to take some other automated action designed to address these types of risks detected during the onboarding process.

While there remain risks related to proper third-party vendor management and governance related to an ML program, which are addressed below, an ML model for ID verification has the potential to reduce risks otherwise presented by human error during the administration of a customer identification procedure and the broader Bank Secrecy Act ("BSA") program, all other considerations remaining equal.

Post-Onboarding Risks

Once an onboarding customer's identification is verified, there remain other risks related to fraud potentially presented during and after the onboarding process.[5] CBI has observed that certain data on the CBI platform serve as indicia of those risks, which include second-party and third-party fraud.[6] For example, those engaged in fraudulent conduct may sometimes change their name to a similar one, or to an alias, to avoid detection and then attempt to set up separate accounts or have wallet addresses under those alternative identities. Similarly, data showing that a single customer is linked to multiple separate accounts and wallet addresses, including ones hosted on other platforms, can be associated with fraudulent activity.

Other common data inputs related to fraud include (i) when a customer buys an asset and immediately sends it to another account or wallet address, or (ii) any unusual transactional activity in a specific wallet address, including anomalous transaction sizes.

ML models can be developed to consume and process this type of information and discern or identify patterns indicative of second- or third-party fraud. The ML models can alert risk managers to conduct additional review and, over time, can learn to automate a

[5] *See* 7 U.S.C. § 6b. *See also* 31 C.F.R. § 1026.220.

[6] Second-party fraud is when a person knowingly gives their identity or personal information to another person, enabling that second person to perform some act to the first person's benefit. Third-party fraud is when a person uses another's identity or personal details without their consent or knowledge in order to gain access to credit or products, commonly referred to as "identity theft."

response such as categorizing a particular account or accounts as “at risk,” imposing a delay on the account’s ability to transmit a transfer, or freezing asset transfers into or out of the account. Deployed in this manner, ML programs can significantly improve the efficiency of reviewing account and transactional information and thereby improve the efficacy of the BSA program.

Summary

Coinbase wishes to stress the following point: the development of AI-based systems is an opportunity for significant innovation and improvement in the financial services sector that should be supported in a responsible manner. These technologies, however, are nascent, and their promise cannot be realized in an overly-prescriptive regulatory environment that cannot keep pace with future developments. The United States has historically been the global hub of innovation and we encourage Treasury to help ensure that remains true.

General Use of AI in Financial Services

Question 1 – *Is the definition of AI used in this RFI appropriate for financial institutions? Should the definition be broader or narrower, given the uses of AI by financial institutions in different contexts? To the extent possible, please provide specific suggestions on the definitions of AI used in this RFI.*

We agree with the scope of the definition adopted in the RFI, particularly in that it includes ML, which is likely the most relevant form of AI for financial institutions today.

Question 2 – *What types of AI models and tools are financial institutions using? To what extent and how do financial institutions expect to use AI in the provision of products and services, risk management, capital markets, internal operations, customer services, regulatory compliance, and marketing?*

As we discuss above in the Introduction, the primary use of ML systems at Coinbase today is for our trade surveillance programs, where we employ ML to improve the identification of manipulative and abusive activity. We are also exploring the expanded use of ML or the new use of other forms of AI to improve the customer experience, including onboarding, fraud detection, customer service, and the use of our website and app.

Question 3 – *To what extent does the type of AI, the development of AI, or AI applied use cases differ within a financial institution? Please describe the various types of AI and their applied use cases within a financial institution. Are there additional use cases for which financial institutions are applying AI or for which financial institutions are exploring the use of AI? Are there any related reputation risk concerns about using AI? If so, please provide specific examples.*

As discussed above, Coinbase's most important current use of ML models is in our trade surveillance programs. These models are trained on large quantities of data and designed to recognize patterns and, importantly, deviations from those patterns, which makes them particularly well-suited to aiding the monitoring efforts of our Surveillance team. One additional feature of ML models that we wish to highlight is the significantly improved degree of transaction insights afforded by their combination with blockchain technology and analytics, which is not available in traditional financial markets.

CBI relies on data from public blockchains in conducting risk management. This data is a compliance enabler because blockchain technology creates a ledger of transactions that are transparent, immutable, and available to any risk managers (as well as to law enforcement or investigation teams). Blockchain-based ledgers are public, distributed, and permanent: anyone can download the ledger and see the entire history of every transaction that has ever occurred on a given blockchain, and no one can change it.[7] This feature allows greater visibility into the counterparties involved in a transaction, and this data can be highly relevant if not necessary to a properly comprehensive review and risk assessment of a customer in the digital asset marketplace.[8]

This additional data facilitates deeper analysis to determine the risk of a specific transaction or asset (an approach known as "know your transaction," or "KYT") instead of relying solely on information and transactions happening within our platform. KYT is groundbreaking for compliance because it is generally immediate (the information is available on the blockchain), independent (it does not have to come from the customer and cannot be tampered with), and dynamic (the risk associated with a customer or transaction can be continually reevaluated based on new blockchain data). This additional, richer dataset available from public blockchains can be continuously processed by ML models to better identify risks—models denied this data would not be able to create the same risk profile of a customer on the platform.[9]

Question 4 – *Are there challenges or barriers to access for small financial institutions seeking to use AI? If so, why are these barriers present? Do these barriers introduce risks for small financial institutions? If so, how do financial institutions expect to mitigate those risks?*

We are not best positioned to answer this question.

[7] *See* Robert Werner *et al.*, *Blockchain Analysis Tool of a Cryptocurrency* 80, 80 (Mar. 2020) https://dl.acm.org/doi/pdf/10.1145/3390566.3391671 ("The blockchain . . . is an immutable ledger, which is stored on a large network of servers worldwide in a decentralized manner. On this ledger, all transactions are stored permanently, transparently and can be accessed by anyone").

[8] *See* Testimony of Grant Rabenn, Director, Financial Crimes Legal at Coinbase, before the U.S. House Committee on Financial Services, *Subcommittee on Digital Assets, Financial Technology, and Inclusion* (Feb. 15, 2024) https://docs.house.gov/meetings/BA/BA21/20240215/116861/HHRG-118-BA21-Wstate-RabennG-20240215.pdf.

[9] KYT also creates an enhanced approach to sanctions compliance in which companies like Coinbase directly screen for crypto addresses identified by the Office of Foreign Assets Control ("OFAC") and can then proactively build out larger networks of high-risk addresses. Before the use of crypto, OFAC was limited to putting static, traditional identifiers—such as names and addresses—on its Specially Designated Nationals List. But with blockchain technology, sanctions compliance can now be based on transactional data, not just personal identifying information. With blockchain analytics, platforms can take ground-truth addresses provided by OFAC to build out and identify much larger networks of high-risk counterparties using blockchain heuristics. They can do this by leveraging immutable transactional data on the blockchain that is unrestricted by private ledgers and can tell them about common ownership.

Actual and Potential Opportunities and Risks Related to Use of AI in Financial Services

Question 5 – *What are the actual and expected benefits from the use of AI to any of the following stakeholders: financial institutions, financial regulators, consumers, researchers, advocacy groups, or others? Please describe specific benefits with supporting data and examples. How has the use of AI provided specific benefits to low-to-moderate income consumers and/or underserved individuals and communities (e.g., communities of color, women, rural, tribal, or disadvantaged communities)? How has AI been used in financial services to improve fair lending and consumer protection, including substantiating information? To what extent does AI improve the ability of financial institutions to comply with fair lending or other consumer protection laws and regulations? Please be as specific as possible, including details about cost savings, increased customer reach, expanded access to financial services, time horizon of savings, or other benefits after deploying AI.*

As we discuss above, greater adoption of AI-enabled technologies like those that we and other market participants are adopting will not only enable registered entities to better meet their own regulatory obligations, but will also help regulators fulfill their missions. In particular, we believe that regulators and self-regulatory organizations that use AI-based systems responsibly for purposes of fraud prevention and deterrence of market manipulation will more effectively ensure orderly markets and investor protection than those that do not. As a consequence, financial markets will enjoy higher integrity and offer greater safety for investors.

We also discuss above the important role that we expect AI-based systems to play in improving the customer experience, in particular by improving fraud detection. We expect that these improvements will be of greater benefit to lower-income consumers, for whom a delay in resolving an instance of fraud may be a more significant hardship. In addition, improvements in onboarding processes and ID-verification may lower the barriers to entry for some consumers who currently transact outside the traditional financial system.

Question 6 – *To what extent are the AI models and tools used by financial institutions developed inhouse, by third-parties, or based on open-source code? What are the benefits and risks of using AI models and tools developed in-house, by third-parties, or based on open-source code? To what extent are a particular financial institution's AI models and tools connected to other financial institutions' models and tools? What are the benefits and risks to financial institutions and consumers when the AI models and tools are interconnected among financial institutions?*

We discuss our use of third party vendors in greater detail below in response to Question 15.

We believe there are significant benefits to AI that can be achieved by better integrating blockchain technology with AI systems.[10] Financial institutions intending to implement AI-based systems must ensure that these systems are secure enough and developed appropriately to meet the expectations of their regulators. It is therefore critically important that a financial institution, or its regulator, be able to verify the data that trains an AI or ML model.

Blockchain can be used to develop solutions that help users and developers ensure that the data and models have not been modified without their knowledge. For example, an API-based service could allow data-owners and AI developers to record time-stamped hashes of datasets and models to ensure their integrity and log the entire process of model development and the datasets used to track the entire lifecycle, in a way that could be made available to third party auditors or regulators. The system could even be directly integrated into ML development tools such as Pytorch. This could help improve the integrity and trustworthiness of models, by making their development process more transparent and secure. It may also be possible to log relevant proofs of the "unlearning" of particular data from models onchain to demonstrate to the satisfaction of regulators that a certain provider's data have been removed from a given model. Logging hashes for data and model outputs onchain can also help to combat deep fakes—for example, applications may be able to ensure the authenticity of the data used by checking digital signatures associated with the source of the data onchain.

Question 7 – *How do financial institutions expect to apply risk management or other frameworks and guidance to the use of AI, and in particular, emerging AI technologies? Please describe the governance structure and risk management frameworks financial institutions expect to apply in connection with the development and deployment of AI. Please provide examples of policies and/or practices, to the extent applicable. What types of testing methods are financial institutions utilizing in connection with the development and deployment of AI models and tools? Please describe the testing purpose and the specific testing methods utilized, to the extent applicable. To what extent are financial institutions evaluating and addressing potential gaps in human capital to ensure that staff can effectively manage the development and validation practices of AI models and tools? What challenges exist for addressing risks related to AI explainability? What methodologies are being deployed to enhance explainability and protect against potential bias risk?*

Coinbase ensures that we have staff / "human" oversight in place to manage risk. For example, we have implemented the following AI-related controls and safeguards throughout our risk management system, including:

[10] Coinbase, *Blockchain for AI*, (Mar. 8, 2024) Blockchain for AI (coinbase.com).

- Employee training pertaining to ethical AI use and best practices;
- Quality assurance testing for all use cases prior to launch to ensure fairness and mitigate unintended bias or discrimination.
- Human in the loop, to ensure fairness and accountability through review of all inputs and outputs to ensure no automated decision making occurs.

Further, our AI team carries out monitoring and testing of its large language solutions and their generated outputs to:

- Detect data or concept drifts that could change its behavior over time ("Model Drift");
- Identify biases or unfairness embedded within the model parameters;
- Ensure model integrity by checking for unexpected changes;
- Benchmark performance against validation datasets; and
- Validate attribution methods.

We make this possible by conducting periodic sampling of the outputs generated by use cases. Human annotators then evaluate the sample for robustness against unintended biases or disclosures before we run validation against a golden test set with ground truthed responses.

Question 8 – *What types of input data are financial institutions using for development of AI models and tools, particularly models and tools relying on emerging AI technologies? Please describe the data governance structure financial institutions expect to apply in confirming the quality and integrity of data. Are financial institutions using "non-traditional" forms of data? If so, what forms of "non-traditional" data are being used? Are financial institutions using alternative forms of data? If so, what forms of alternative data are being used?*

As we discuss above, ML models can be trained on existing trading data to better detect instances of manipulation and abusive behavior. Blockchain technology deepens the available dataset for training these models, as the blockchain's entire transaction history can be used for training. We also believe, as discussed above, that blockchain technology offers an important opportunity to improve the data governance of these training sets, as the data that is incorporated can itself be immutably recorded in a transparent blockchain.

Fair Lending, Data Privacy, Fraud, Illicit Finance, and Insurance

Question 9 – *How are financial institutions evaluating and addressing any increase in risks and harms to impacted entities in using emerging AI technologies? What are the specific risks to consumers and other stakeholder groups, including low- to moderate-income consumers and/or underserved individuals and communities (e.g., communities of color, women, rural, tribal, or disadvantaged communities)? How are financial institutions protecting against issues such as dark patterns – user interface designs that can potentially manipulate impacted entities in decision-making – and predatory targeting emerging in the design of AI? Please describe specific risks and provide examples with supporting data.*

As we discuss above in response to Question 7, we incorporate a significant degree of staff oversight and testing in our ML models to mitigate potential risks. We believe that many of the risks associated with a poorly trained AI or ML model that this question describes can be mitigated by a transparent training set, which can be audited to assess for the fairness of the data incorporated into the model. In general, as we discuss above in the Introduction and in response to Question 5, we believe that the additional incorporation of AI and ML models in financial services, when conducted responsibly, is likely to improve the provision of financial services for all consumers.

Question 10 – *How are financial institutions addressing any increase in fair lending and other consumer-related risks, including identifying and addressing possible discrimination, related to the use of AI, particularly emerging AI technologies? What governance approaches throughout the development, validation, implementation, and deployment phases do financial institutions expect to establish to ensure compliance with fair lending and other consumer-related laws for AI models and tools prior to deployment and application? In what ways could existing fair lending requirements be strengthened or expanded to include fair access to other financial services outside of lending, such as access to bank accounts, given the rapid development of emerging AI technologies? How are consumer protection requirements outside of fair lending, such as prohibitions on unfair, deceptive and abusive acts and practices, considered during the development and use of AI? How are related risks expected to be mitigated by financial institutions using AI?*

We are not best positioned to discuss the impact on fair lending.

With respect to consumer protection generally, as we discuss above in the Introduction, we believe there are significant opportunities to mitigate, identify, and resolve instances of consumer fraud or other scams using ML models. These models can improve customer ID verification to reduce instances of identity theft and enable financial institutions to better identify potentially fraudulent activity.

Question 11 – *How are financial institutions addressing any increase in data privacy risk related to the use of AI models, particularly emerging AI technologies? Please provide examples of how financial institutions have assessed data privacy risk in their use of AI. In what ways could existing data privacy protections (such as those in the Gramm-Leach-Bliley Act (Pub. L. No. 106-102)) be strengthened for impacted entities, given the rapid development of emerging AI technologies, and what examples can you provide of the impact of AI usage on data privacy protections? How have technology companies or third-party providers of AI assessed the categories of data used in AI models and tools within the context of data privacy protections?*

Coinbase places a strong emphasis on the importance of protecting our customers' data throughout our organization. In particular with respect to AI and ML models, we have implemented the following controls and safeguards throughout our risk management system, including:

- Transparency in our privacy policy and disclosures ensuring users/customers are notified when interacting directly with a chatbot powered by a large language model ("LLM").
- Data minimization and access controls to limit data input into the models to ensure improper processing and or sharing of information outside the scope of what the user has permitted.
- Zero retention architecture to minimize unauthorized dissemination or training on users' data.

We have established a unique AI risk & control framework which has been embedded into our Security & Privacy shared services functions, including security assessments and privacy assessments. This framework promotes consistent documentation, tracking, and treatment of identified risks through a centralized Security & Privacy risk register, as well as through centralized risk management.

Question 12 – *How are financial institutions, technology companies, or third-party service providers addressing and mitigating potential fraud risks caused by AI technologies? What challenges do organizations face in countering these fraud risks? Given AI's ability to mimic biometrics (such as a photos/video of a customer or the customer's voice) what methods do financial institutions plan to use to protect against this type of fraud (e.g., multifactor authentication)?*

As we discuss in the Introduction, there are significant customer fraud benefits to the adoption of ML models in onboarding and post-onboarding processes. These models assist with both customer ID verification and the identification of fraudulent activity conducted through a customer's account.

Question 13 – *How do financial institutions, technology companies, or third-party service providers expect to use AI to address and mitigate illicit finance risks? What challenges do organizations face in adopting AI to counter illicit finance risks? How do financial institutions use AI to comply with applicable AML/CFT requirements? What risks may such uses create?*

We address this question in greater detail in the Introduction and in our response to Question 3. We believe that AI and ML models will become an important tool in countering financial crime risks.

Question 14 – *As states adopt the NAIC's Model Bulletin on the Use of Artificial Intelligence Systems by Insurers and other states develop their own regulations or guidance, what changes have insurers implemented and what changes might they implement to comply or be consistent with these laws and regulatory guidance? How do insurers using AI make certain that their underwriting, rating, and pricing practices and outcomes are consistent with applicable laws addressing unfair discrimination? How are insurers currently covering AI-related risks in existing policies? Are the coverage, rates, or availability of insurance for financial institutions changing due to AI risks? Are insurers including exclusions for AI-related risks or adjusting policy wording for AI risks?*

We are not best positioned to answer this question.

Third-Party Risks

Question 15 – *To the extent financial institutions are relying on third-parties to develop, deploy, or test the use of AI, and in particular, emerging AI technologies, how do financial institutions expect to manage third-party risks? How are financial institutions applying third-party risk management frameworks to the use of AI? What challenges exist to mitigating third-party risks related to AI, and in particular, emerging AI technologies, for financial institutions? How have these challenges varied or affected the use of AI across financial institutions of various sizes and complexity?*

In addition to developing its own AI models, Coinbase partners with third-party vendors that use ML models in limited ways to deliver their products, specifically to CDE. In selecting these partners, Coinbase has not sought particular AI expertise or models, but rather the best product solutions, which might happen to leverage AI within the product-solution scope.

Coinbase has followed best practices related to the risk management of such vendors, including engaging in appropriate due diligence of potential partners before the relationships begin. This review involves assessing whether a potential vendor could satisfy Coinbase's own policies and procedures where necessary, comply with applicable regulations, protect any data required to be shared with the vendor, and effectively allow monitoring of the vendor once it begins providing the product or service.

Although the principles of third-party risk management remain the same for vendors leveraging AI or ML tools, they do present some novel circumstances for managing third-party risks. In the realm of information and regulatory systems, a vendor's products tend to be powered by rules-based, algorithmic software where the processing of certain types of data will lead to predictable results. By and large, anomalies in output will be the result of anomalies in the data (unless, of course, there is a processing malfunction). With AI systems, data inputs feed into evolving decision-making paradigms—changes in data can lead to changes in the methods of processing itself.

Thus, due diligence and ongoing management of a vendor using AI require a level of understanding of the AI or ML tool itself. Achieving this understanding would involve a rigorous assessment to confirm that the correct data inputs are being ingested and turned into appropriate decision-making paradigms by the AI system. This, of course, requires adequate transparency into the design and functioning of these tools, which must be present at all stages of the vendor relationship.

Assessing whether a vendor using AI would enable Coinbase to meet any and all of its own relevant regulatory obligations requires an even deeper level of scrutiny and governance. Depending on the specific purpose of the vendor product, features of this review might include rigorous review of contractual terms; business leader accountability of and sign off for the vendor relationship; an independent evaluation of the vendor's infrastructure, controls, risks, and effectiveness of their controls (i.e., a SOC report); a review of any licensure and insurance for the vendor; and cybersecurity assessments.

In sum, an adequate level of technical understanding of these systems by the end user, combined with relatively greater transparency into the systems themselves provided by the vendor, should be the hallmarks of appropriate risk management of these relationships. This transparency should also include the vendor providing necessary access to data and information when necessary to review and deconstruct operational incidents. One necessary control that should be in place to achieve this goal is including a regular audit of the vendor and its AI systems as a feature of the vendor contract.

Question 16 – *What specific concerns over data confidentiality does the use of third-party AI providers create? What additional enhancements to existing processes do financial institutions expect to make in conducting due diligence prior to using a third-party provider of AI technologies? What additional enhancements to existing processes do financial institutions expect to make in monitoring an ongoing third-party relationship, given the advances in AI technologies? How do financial institutions manage supply chain risks related to AI?*

As discussed in response to Question 15, Coinbase has implemented a comprehensive risk management strategy to address risks arising from the use of third-party models or vendors. This strategy includes several key components with respect to third-party AI usage:

- Enhancements to our third party risk management processes to include explicit checks for third-party use of AI. Additionally, we have embedded unique AI control requirements into our third party security assessments.
- Requirement that any third-party hosted LLM that Coinbase obtains a license to use has been configured for zero-retention of data ("incognito").
- Ensuring that third party LLM providers do not retain any Coinbase data, knowledge bases, or user inputs processed by the models.
- Ensuring that third party staff do not have access to Coinbase data, knowledge bases, or user inputs processed by the models.
- Ensuring that logs collected by the third party LLM provider, as part of their ongoing management of the model, do not contain Coinbase data or inputs processed by the models.

Question 17 – *How are financial institutions applying operational risk management frameworks to the use of AI? What, if any, emerging risks have not been addressed in financial institutions' existing operational risk management frameworks? How are* *financial institutions ensuring their operations are resilient to disruptions in the integrity, availability, and use of AI? Are financial institutions using AI to preserve continuity of other core functions? If so, please provide examples.*

We are not best positioned to answer this question.

Further Actions

Question 18 – *What actions are necessary to promote responsible innovation and competition with respect to the use of AI in financial services? What actions do you recommend Treasury take, and what actions do you recommend others take? What, if any, further actions are needed to protect impacted entities, including consumers, from potential risks and harms? Please provide specific feedback on legislative, regulatory, or supervisory enhancements related to the use of AI that would promote a financial system that delivers inclusive and equitable access to financial services that meet the needs of consumers and businesses, while maintaining stability and integrity, protecting critical financial sector infrastructure, and combating illicit finance and national security threats. What enhancements, if any, do you recommend be made to existing governance structures, oversight requirements, or risk management practices as they relate to the use of AI, and in particular, emerging AI technologies?*

As we discuss in the Introduction, this is a nascent and rapidly developing technology whose promise and risks are best addressed in a minimally prescriptive manner with guidance from regulators, but without prescriptive requirements that could risk the potential benefits of the technology. We encourage Treasury and financial services regulators to partner with those in the financial services industry who are working on these developments, to make sure that regulatory expectations and development goals work constructively.

Question 19 – *To what extent do differences in jurisdictional approaches inside and outside the United States pose concerns for the management of AI-related risks on an enterprise-wide basis? To what extent do such differences have an impact on the development of products, competition, or other commercial matters? To what extent do such differences have an impact on consumer protection or availability of services?*

Crypto markets, more so than other financial markets, are truly global. As a consequence, crypto markets require a significant degree of cross-border regulatory cooperation and harmonization. This applies to rules surrounding the use of AI-based systems as well. For example, to the extent we use ML models to monitor trading on CBI, it is not clear how more restrictive rules on the use of AI in one jurisdiction would impact our global monitoring program. In the same vein, we recommend learning from existing regulations to enhance any US approach, ensuring the same uncertainties and risks are not replicated. It is crucial that any new regulations provide clear and precise guidelines to avoid unintended negative consequences on the global market. By doing so, innovation and operational efficiency is fostered, while maintaining robust oversight.

We would also be opposed to localization requirements in this area, as we have in other areas of financial regulation. For example, a rule requiring that an AI or ML model trained on data from one country be used *only* in that country would limit the effectiveness of the AI or ML model in identifying manipulative trading behavior and thus ultimately harm consumers—the same way that requiring customer crypto assets to be custodied in one country makes those assets more vulnerable than if they can be custodied on a distributed global basis. Global distribution improves the quality of consumer protection and helps ensure that services are always available in the event of local disruption. But global distribution is made more difficult by global regulatory variation.

Conclusion

Coinbase commends Treasury for releasing the RFI and for appreciating the impact that machine learning technology will have, and is already having, on financial markets and the financial services industry. Indeed, Coinbase believes that ML methods will be an important tool for growth, innovation, and consumer protection in the financial services sector. Similarly, crypto and blockchain technology can deliver a more fair, accessible, efficient, and transparent system to transfer value and ownership. Together, the dual emerging technologies of AI and digital assets have the potential to transform multiple industries, with AI addressing large-scale problem solving, and digital asset innovations ensuring the authenticity and provenance of underlying information.

Specifically, and as noted above, crypto brings with it tools—namely, an immutable, public ledger—that is unavailable in traditional finance. Investigations on a blockchain are easier. The market surveillance methods enabled by machine learning paired with blockchain technology have the potential to create a more secure system than the current traditional finance approach. Well-designed regulation of both AI and digital assets will provide the market the certainty and workability it needs to power these innovations.