coinbase

To:
Directorate-General for Financial Stability, Financial Services and Capital Markets Union
European Commission
1049 Bruxelles/Brussel
Belgium

13 September 2024

Re: Targeted Consultation on Artificial Intelligence in the Financial Sector

Coinbase Global, Inc. (together with its subsidiaries, **Coinbase**) appreciates the opportunity to respond to the targeted consultation (the **Consultation**) published by the European Commission's Directorate-General for Financial Stability, Financial Services and Capital Markets Union (the **Commission**).

Coinbase started in 2012 with the idea that anyone, anywhere, should be able to send and receive Bitcoin easily and securely. Today, we are working to responsibly adopt AI and machine learning technologies in our mission to expand economic freedom. 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 embracing the use of AI. In particular, we believe that the use of AI-based systems to prevent fraud and deter market manipulation and other potential abuses can more effectively ensure orderly markets and investor protection. But the adoption of AI must be done responsibly. We encourage the Commission and financial services regulators to partner with financial institutions to make sure that appropriate protections are in place while not stifling innovation.

Coinbase appreciates the Commission'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

Tom Duff Gordon, Vice President, International Policy, Coinbase

Introduction

Coinbase applauds the release of the Consultation because it demonstrates concretely that the Commission 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 (**GenAI**)[2] products and processes is associated with supervision of market activities in Coinbase, Inc. (**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. 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. Laws that create complex or potentially unclear rules for the use of AI 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 used by the Consultation,[3] is used by Coinbase in certain carefully governed ways to improve processes within the company, with promising results. Beyond ML, systems using GenAI are not yet widely deployed by Coinbase for product development and production or risk management.

---

[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] Generative AI is a type of artificial intelligence (**AI**) that can create new content, such as text, images, music, or code.

[3] *See* Consultation at 4 ("any machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment and 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"). We use AI throughout this response as a general term to refer to both ML and GenAI.

It is important to note that the relationship between AI and financial services, particularly services using blockchain technology, is a reciprocal and mutually beneficial one. Innovations in one can be used to advance the other. We discuss at length in our response the ways in which AI is helping the financial services sector improve, but we also wish to note that innovation in blockchain-based financial services can improve the quality and utility of AI. In August, Coinbase witnessed the first AI-to-AI blockchain-based transaction on our blockchain, Base.[4] This transformative technology will empower AI tools to send and receive funds by accessing on-chain wallets, allowing them to complete transactions independently and provide much greater utility for users overall.

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.[5] 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.

---

[4] Brian Armstrong, *AIs are now paying other AIs with crypto*, (August 30, 2024) Brian Armstrong on X: "AIs are now paying other AIs with crypto" / X.

[5] 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; Coinbase, *Re: Request for Information on Uses, Opportunities, efficiand Risks of Artificial Intelligence in the Financial Services Sector* (Aug. 12, 2024), *available at* https://assets.ctfassets.net/o10es7wu5gm1/6uO6LHCnY5Rm7fYeiM3R6P/f99c7948d2445605917348e10e9b7422/2024-08-12-UST_RFI_on_Artificial_Intelligence.pdf.

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. AI-enabled technologies stand 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.

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 GenAI systems to improve the overall customer experience. GenAI and ML models hold great promise in improving the customer onboarding process, including in particular the preliminary steps in verifying customer identification, such as address verification. (Coinbase does not use AI models to analyze biometric data, and does not plan to do so.) 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, AI-enabled technologies can assist with the ID-verification process and mitigate first-party fraud. CBI has leveraged third-party ML models to further automate the onboarding process and reduce the risks of human error that might enable fraudulent behavior. These models 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 AI program, which are addressed below, an AI 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 anti-money laundering 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. 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.

[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."

AI-enabled technologies can be developed to consume and process this type of information and discern or identify patterns indicative of second- or third-party fraud. The AI models can alert risk managers to conduct additional review and, over time, can learn to automate a 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, AI-enabled technologies can significantly improve the efficiency of reviewing account and transactional information.

Enhancing Customer Recommendations

Finally, Coinbase is using ML to enhance the customer experience by providing more personalized notifications and recommendations. ML models can help provide more curated information to customers when they visit the Coinbase website or app, and can enhance customer service by crafting customized responses to inquiries. Coinbase continues to explore ways in which to use AI-enabled technologies to improve customer experience.

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.

Part 1: GENERAL QUESTIONS ON AI APPLICATIONS IN FINANCIAL SERVICES

1.1. Use of AI

Question 1. Are you using or planning to use AI systems?

- Yes, we are already using AI systems.

Question 2. What are the positive things you encounter when using AI?

Open answer/Please explain and give examples when possible:

As discussed in the Introduction, 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 is 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

---

[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.

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]

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.

In addition, as we discuss in the Introduction, financial services and AI can exist in a positive, reciprocal relationship where each serves to enhance the other. Our experiences enabling AI-to-AI crypto transactions are an early step in using financial innovation to improve the utility of AI systems.

Question 3. What are the negative things you encounter when using AI?

Open answer/Please explain and give examples when possible:

Incorporating AI-based systems responsibly requires additional layers of risk management. If not approached properly, there is a risk of misusing data or AI systems leading to unfavorable outcomes. Coinbase ensures that we have staff / "human"

[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.

oversight in place to manage risk. For example, we have implemented the following AI-related controls and safeguards throughout our risk management system, including:

- 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 inputs and outputs to ensure no automated decision making occurs.

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

- Detect data or concept drifts that could change their behavior over time ("Model Drift");
- Identify biases or unfairness;
- 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 4. Will you be deploying AI for new or additional processes within your organization?

- Yes, which ones?

In addition to the current and planned uses we discuss above, Coinbase is investigating additional uses for AI-based systems. For example, methods currently used by CBI (*e.g.*, to close out false positive fraud alerts) 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. AI-enabled technologies stand to provide a new perspective into how CDE's markets and participants operate, as well as potentially identifying and highlighting new disruptive practices.

We have also recently discussed our efforts to use ML to predict traffic and scale databases.[10] Crypto markets can be volatile. On the Coinbase platform, spikes of user activity and traffic can occur suddenly and quickly, and then disappear just as fast. We handle those changing traffic patterns and workloads by scaling up and adding resources

[10] Coinbase, *How Coinbase is Using Machine Learning to Predict Traffic and Scale Databases*, (Aug. 26, 2024) How Coinbase is Using Machine Learning to Predict Traffic and Scale Databases.

in times of high traffic. Scaling up is not an instant process, however, and beginning to scale when traffic is already high is often too late. Therefore, we have developed an automatic scaling solution that uses ML to predict the traffic spikes and trigger a scale up before the traffic arrives.

Question 5. Are you developing or planning to develop in-house AI applications?

- Yes, please explain.

Coinbase primarily develops its own ML models to assist with risk assessment and customer recommendations. We also partner with third-party vendors that use GenAI models in limited ways to deliver their products, specifically to CDE. In selecting these partners, Coinbase has sought the best product solutions, which might happen to leverage AI within the product-solution scope.

Question 6. Which tools are you using to develop your AI applications?

Examples: machine learning, neural networks, natural language processing, large language models, etc.

Open answer/Please explain and give examples when possible:

As we discuss above, Coinbase's primary use of AI-based systems today is with machine learning.

1.2. Benefits of using AI applications in financial services

Question 7. Please score the following benefits from most significant (10) to least significant (1):

- 10 -- Fraud detection: AI algorithms can analyze large amounts of data to detect patterns and anomalies that may indicate fraudulent activity, helping to reduce financial losses for businesses and customers.
- N/A -- Risk management: AI can analyze and predict market trends, assess credit risks, and identify potential investment opportunities, helping financial institutions make more informed decisions and manage risks more effectively.
- 8 -- Automation of routine tasks: AI can automate repetitive tasks such as data entry, transaction processing, and document verification, freeing up time for employees to focus on more complex and strategic activities.
- 8 -- Cost savings: by automating processes and improving efficiency, AI can help financial institutions reduce operational costs.

- N/A -- Personalized financial advice: AI can analyze customer data to provide personalized financial advice and recommendations, helping customers make better financial decisions and improve their financial well-being.
- N/A -- Compliance and regulatory support: AI can help financial institutions stay compliant with regulations by analyzing and interpreting complex regulatory requirements and monitoring transactions for suspicious activities.
- N/A -- Enhanced decision-making: AI can analyze large amounts of data and provide insights that can help financial institutions make better investment decisions, assess credit risks, and optimize their operations.
- N/A -- Improved security: AI can enhance security measures by identifying potential security threats, detecting unusual patterns of behavior, and providing real-time alerts to prevent security breaches.
- N/A -- Streamlined processes: AI can streamline various financial processes, such as loan underwriting, account opening, and claims processing, leading to faster and more efficient services for customers.
- 8 -- Improved customer service: AI can be used to provide personalized and efficient customer service, such as chatbots that can answer customer queries and provide assistance 24/7.

Question 8. What are the main benefits/advantages you see in the development of your AI applications?

Open answer/Please explain and give examples when possible:

We address the benefits of AI more generally in our response to Question 2. We believe it is also important to note, however, the significant benefits to AI that can be achieved by better integrating blockchain technology with AI systems.[11] 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 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

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

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.

1.3. Challenges and risks when using AI applications in financial services

Question 9. Please score the following challenges and risks from most significant (10) to least significant (1):

Coinbase is not best positioned to assess risk related to bias and reputational risks; therefore, we have responded "N/A." As discussed throughout our response, we do not view data access as problematic given that data availability is another benefit of our efforts to integrate blockchain technology into AI-based systems. We are also implementing policies on transparency and security that help mitigate any risks in those areas. Finally, while we see enormous potential for innovation and the ability to leverage AI in combination with crypto and blockchain technologies, we feel the risks in this area are well managed.

- Lack of access to the required data, in general. **N/A**
- Lack of access to the data in an appropriate digital format. **N/A**
- Lack of access to appropriate data processing technology, e.g. cloud computing.
- Data privacy: it is crucial to ensure that sensitive financial information remains confidential. **N/A**
- Lack of trust in relation to performance levels/ security aspects/ certified solutions/ reliability of the technology. **N/A**
- Regulatory compliance with financial regulation: financial services are heavily regulated and not all types of AI applications are in line with requirements under these regulations. **N/A**
- Innovation: the ability to leverage on combining AI with other technologies to enhance its potential and generate new services? **N/A**
- Transparency and explainability: AI algorithms can be complex and opaque. It can be difficult for humans to understand how AI arrives at certain conclusions, which can create issues of trust and accountability. **N/A**
- Bias and discrimination: AI models are trained using data, and if the data is biased, the AI model can also be biased, leading to unfair outcomes. **N/A**
- Reputational risk from undesirable AI behavior or output. **N/A**
- Liability risks: legal uncertainty on who bears the liability in case of damages generated by the malfunctioning of the AI applications. **N/A**

- Skills gap: the development of AI requires specific tech skills, and there is a shortage of such skills. **N/A**
- Dependability: as financial institutions rely more and more on AI; the dependability of these systems becomes paramount. Any malfunction or error (e.g. in risk management) can lead to significant financial losses. **N/A**
- Job displacement: the use of AI can potentially automate certain roles in the financial sector leading to job displacement. **N/A**
- Cybersecurity: AI systems could be targeted by cybercriminals, leading to potential data breaches or manipulation of AI systems. **N/A**
- Integration challenges: integrating AI technologies with existing systems and processes can be complex and expensive. **N/A**
- Additional cost: the deployment and use of AI requires up-front investment and ongoing resources (acquiring or developing applications, keeping them up to date, training/skills). **N/A**

Question 10. What are the main difficulties/obstacles you are facing in the development of your AI applications?

Open answer/Please explain and give examples when possible:

Coinbase faces challenges common to development of all AI applications and ML models, including issues such as latency, throughput, and developing guardrails around potential mistakes.

Question 11. Please rank the potential negative impact that widespread use of AI can have on the following risks. 8 being the highest risk.

- N/A -- Operational risks
- N/A -- Market risks
- N/A -- Liquidity risks
- N/A -- Financial stability risks
- N/A -- Market integrity risks
- N/A -- Investor protection risk
- N/A -- Consumer protection risk
- N/A -- Reputational risk

Please explain your answer to the previous question and give examples when possible:

Although we believe in the promise of AI and are implementing ML technologies where appropriate, we are not in a position to accurately speculate about the risks of widespread use of AI. Because adoption of these technologies remains nascent, any response would be hypothetical and likely inaccurate. As a technology company, Coinbase has first-hand experience with the rapid pace at which technology develops.

Where technology poses risks, mitigating those risks is often a prerequisite to widespread adoption. Moreover, while we believe that AI has the potential to improve financial services in many dimensions, it is not clear to us that AI use will be widespread across all dimensions.

Question 12. AI may affect the type and degree of dependencies in financial markets in certain circumstances, especially where a high number of financial entities rely on a relatively small number of third-party providers of AI systems. Do you see a risk of market concentration and/or herding behavior in AI used for financial services?

- No, please explain.

Coinbase is not best positioned to answer this question.

1.4. AI and compliance burden

Question 13. Can AI help to reduce the reporting burden?

- Yes, in which areas do you see AI reducing reporting burden?

As we discuss in response to Question 2, we believe that AI-based systems, especially in combination with blockchain technology, serve as a compliance enabler. Our experience is that ML systems improve the quality of fraud detection and customer onboarding in ways that make regulatory reporting faster and more accurate. We expect these trends to continue as our models are likely to get better at detecting fraudulent activity the more data can be used to train them.

Question 14. Do you think AI can facilitate compliance with multiple regulatory standards across the EU and thus facilitate market integration or regulatory compliance? For example, would you consider it feasible to use AI for converting accounting and financial statements developed under one standard (e.g. local GAAP) to another standard (e.g. IFRS)? Please elaborate.

Open answer/Please explain and give examples when possible:

As we discuss above, we believe that AI-enabled technologies can improve regulatory reporting for financial institutions; this includes reporting to comply with varying standards. It is important to note, however, that using AI-based systems in novel ways such as contemplated by this question will require financial regulators across jurisdictions to facilitate rather than stifle innovation in the use of AI. For example, it may be possible for a financial institution to develop a degree of automated compliance reporting that takes into account different jurisdictions' obligations. If, however, one jurisdiction makes it

impossible for financial institutions to test such a system, for example by imposing onerous requirements on the use of AI, that could stop the overall improvement in compliance reporting across jurisdictions.

1.5. Data access

Question 15. In order to develop AI applications, do you need access to external datasets that you currently don't have access to?

- No

Question 16. Which datasets would you need to develop meaningful AI applications and for which purpose / use case?

Open answer/Please explain and give examples when possible:

Coinbase is not best positioned to answer this question.

Question 17. Do you face hurdles in getting access to the data you need to develop AI applications in financial services?

- No

As we discuss above, the availability of data is another benefit of integrating blockchain technology in AI-based systems. Data on a blockchain is stored publicly, so that any user can download the entire transaction history of the chain. This allows any user, including financial institutions themselves, to build models based on publicly available and transparent data sets.

Question 18. Are you familiar with the EU Data Hub, a data sharing tool for supervisors and financial companies?

- Yes, do you think it can improve access to data?

We recognize that the ESMA data hub provides data that can be useful for market participants. In particular, the data hub can be a useful tool for the further development of AI-based regulatory and supervisory technology solutions for risk and compliance management.

Question 19. Should public policy measures (e.g. legislative or non-legislative) encourage the exchange of data between market participants, which can be used to train AI systems for use cases in finance?

- Yes. Which type of measures do you propose?

In general, we support a market-led approach to open finance that prioritises use cases with clear customer demand. We also encourage public policy makers to support the adoption of blockchain technology, which, by creating a transparent ledger of publicly available data, resolves the need for specific data sharing arrangements between market participants.

1.6. Business model

Question 20. Has AI changed your business model?

- No

Question 21. Which parts of the value chain are being improved with AI?

Open answer/Please explain and give examples when possible.:

As we discuss above, 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 22. Are there functions that cannot/would not be improved by AI?

Open answer/Please explain and give examples when possible:

We are not best positioned to answer this question.

1.7. General purpose AI

Question 23. Do you use general purpose AI models, including generative AI, and their respective reference architectures?

- Yes, please explain why you want to opt for these AI models in your organization.

We support the use of certain general purpose AI models because we believe they can offer substantial improvements in internal efficiency and the customer experience.

Question 24. How do you plan to operationalise and adopt general purpose AI at scale?

Open answer/Please explain and give examples when possible:

Our primary plans for the adoption of general purpose AI include the use of generally available large language models (**LLMs**). These tools improve our internal efficiency and

have the ability to enhance employee productivity. We also currently use and plan to make use of general purpose AI for customer support (*e.g.*, generating answers to customer inquiries).

Question 25. How does the increasing availability of general purpose AI models, including generative AI applications, impact the need to access new datasets?

Open answer/Please explain and give examples when possible:

As we discuss above, we believe that integrating blockchain technology with AI models poses a significant opportunity to the quality of the AI models regardless of the type of model. Blockchains contain the entire transaction history for each chain and so are able to provide a level of data quality and completeness that has not previously been possible.

Question 26. Compared to traditional AI systems such as supervised machine learning systems, what additional opportunities and risks are brought by general purpose AI models?

Open answer/Please explain and give examples when possible:

As we discuss in other responses, general purpose AI models can offer substantial speed and efficiency gains in areas such as research and development, coding, risk assessment, and customer support. The challenges associated with such models include latency, throughput, and ensuring appropriate guardrails.

Question 27. In which areas of the financial services value chain do you think general purpose AI could have a greater potential in the short, medium and long term?

Open answer/Please explain and give examples when possible:

Customer experience is one area that could benefit from general purpose AI. Coinbase is in the process of deploying existing LLM solutions that will be used to improve the customer experience. We approach changes to our customer relationships cautiously and responsibly, including implementing requisite guardrails prior to LLM usage, and so expect to see significant improvements over a relatively longer timeframe.

1.8. AI Governance in relation to non-high risk use cases, and which are not subject to specific requirements under the AI Act

Question 28. Have you developed, or are you planning to develop an AI strategy or other relevant guidelines within your organization for the use of AI systems?

- Yes, which ones?

As we discuss in response to Question 3, we have developed a strong risk management program related to AI that includes staff oversight at every level.

In addition, Coinbase places a strong emphasis on the importance of protecting our customers' data throughout our organization. In particular with respect to general purpose AI 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 an 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 of the underlying LLMs on users' data.

We have established a unique AI risk and control framework that has been embedded into our Security and 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 and Privacy risk register, as well as through centralized risk management.

Question 29. Have you put in place or are you planning to put in place governance and risk management measures to ensure a responsible and trustworthy use of AI within your organization?

- Yes, which ones?

We address this question more fully in response to Question 28.

1.9. Forecasts

Question 30. What are the main evolutions to be expected in AI in finance?

Open answer/Please explain and give examples when possible:

We believe that AI-enabled technologies will become increasingly used by the financial services sector and present an important opportunity to strengthen both the operations and the compliance programs of financial institutions. As discussed in our Introduction, the use of data from public blockchains acts as a compliance enabler by creating a ledger of transactions that are transparent, immutable, and available to any risk managers (as well as to law enforcement or investigation teams). 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.[12]

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. 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.[13]

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.

Question 31. Which financial services do you expect to be the most impacted by AI?

Open answer/Please explain and give examples when possible:

As we have discussed at several points above, we believe that blockchain technology can significantly improve the training and reliability of AI models. As a result, we expect that AI will have a greater positive impact on financial institutions that use blockchain technology. We also anticipate that crypto and blockchain technology will help AI tools become much more effective by facilitating AI-to-AI blockchain-based payments. Although AI tools cannot hold bank accounts or access traditional payment rails, crypto and blockchain can enable them to send and receive funds, complete transactions independently, and provide much greater utility for users overall.

[12] *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.

[13] 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.

Question 32. Do you have any additional information to share?

Coinbase believes that AI-enabled technologies 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 are 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.

Part 2: QUESTIONS RELATED TO SPECIFIC USE CASES IN FINANCIAL SERVICES

Question 33. In which sector are you using AI? You may select more than one answer.

- Banking and payments
- Market infrastructure
- Asset management

Banking and payments (if selected)

Question BANKING 1. For which use case(s) are you using/considering using AI?

Coinbase uses ML in similar ways across various use cases that encompass both payments and trading. Because most of our ML use relates to market infrastructure, we have responded to questions in that section.

Market infrastructure (if selected)

Question MARKET 1. For which use case(s) are you using/considering using AI?

Open answer. Examples: risk management, sustainable finance, regulatory compliance, etc:

As we discuss above, we are primarily using ML with respect to our trade surveillance and fraud detection functions. There are, however, clear extensions of these current use cases to AML and other compliance obligations. We are also actively working to incorporate AI-based systems to improve the customer experience, from onboarding, to fraud detection, to customer service.

Question MARKET 2. What are the opportunities that AI brings to your use case?

Open answer/Please explain and give examples when possible.

As we discuss in the Introduction, ML technology has enabled us to determine specific instances of fraud or manipulative activity with much greater accuracy. ML models are designed to absorb and analyze large amounts of data to spot patterns and deviations from those patterns. That ability makes these models a perfect tool for use as part of our overall transaction monitoring program, which must sort through and analyze all trading data that occurs on our system.

Question MARKET 3. What are the main challenges and risks that AI brings to your use case (e.g. discrimination, opacity of the AI application developed, difficult to control/supervise it, etc.)?

Open answer/Please explain and give examples when possible.

As we discuss above, it is important for financial institutions to incorporate AI models responsibly, which requires the adoption of specific risk management programs and controls for customer privacy.

Question MARKET 4. What is the main barrier to developing AI in your use case (e.g. lack of skills and resources, readiness of the technology, high regulatory costs for compliance with the relevant frameworks, etc.)?

Open answer/Please explain and give examples when possible.

We do not have any significant barriers to benefiting from ML tools. As discussed above, we believe that regulators have an opportunity to protect consumers and enhance compliance by encouraging the widespread adoption of KYT and similar approaches that incorporate the benefits of blockchain technology into ML.

Question MARKET 5. Does AI reduce or rather increase bias and discrimination in your use case?

Please explain and give examples when possible.

As we discuss above, we expect AI-enabled technologies to serve an important role in improving the customer experience, in particular by improving fraud detection and the onboarding process. We believe that each of these improvements will likely be of greater benefit to lower-income consumers for whom delays in resolving fraudulent activity may be of more significant hardship and who may face barriers to entry in the traditional onboarding process.

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.

Question MARKET 6. Has general purpose AI opened new possibilities or risks in your use case?

- No

Question MARKET 7. On whom do you rely for the development of your AI solutions?

- Partial collaboration with external providers

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 sought the best product solutions, which may 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.

Coinbase has implemented a comprehensive risk management strategy to address risks arising from the use of third-party LLM 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.

Asset management (if selected)

Question ASSET MANAGEMENT 1. For which use case(s) are you using/considering using AI?

Open answer. Examples: risk management, individual and collective portfolio management, regulatory compliance, trades monitoring, robo-advice, customer service, sustainable finance, etc.

We use AI-based tools to assist in our asset management research and development processes. These tools are particularly beneficial for accelerating improvements to our code base and algorithmic infrastructure. They also inform our research efforts as we evaluate new sources of information, new signals, and new approaches to portfolio construction for potential inclusion in our processes.

Question ASSET MANAGEMENT 2. What are the opportunities that AI brings to your use case?

AI-based tools improve the speed and efficiency of the R&D process for asset management.

Imagine a researcher has an idea for a new investment signal. AI tools allow for better information search and summarization, which can help the researcher to:

- Identify what data source(s) are available for the signal;
- Provide suggested code for transforming the underlying data into a probabilistic model, in the coding language of choice;
- Identify which package dependencies and libraries are necessary or useful to support the code; and
- Annotate code with plain-language descriptions of the actions.

By reducing the time required to navigate these steps in the R&D process, researchers and developers can get to better business outcomes in shorter timelines.

Question ASSET MANAGEMENT 3. What are the main challenges and risks that AI brings to your use case (e.g. discrimination, opacity of the AI application developed, difficult to control/supervise it, etc.)?

When incorporating AI in asset management, it can be challenging to preserve the economic or theoretical interpretation behind investment strategies. AI models may identify relationships in data and between variables that cannot be readily explained. While this is not inherently problematic, it is a distinction relative to intuitive investment models.

One of the primary concerns with using AI in asset management relates to information leakage. Asset managers must act with care to maintain discretion around their trading and investing behavior. New models under development provide potential sources of investment returns. While AI can improve the development timelines, care must be taken to preserve information security. Confidentiality is key to maximizing investor outcomes and minimizing risks. Ensuring that asset management data and/or proprietary infrastructure are not allowed to be passed on to AI models as training data is of primary importance.

While AI tools, code, and text suggestions are helpful in our processes, they cannot be used without human review and evaluation. Code suggestions can be erroneous at times, and/or may lag behind the latest developments in software and package releases. For complicated tasks, the suggested logic may not be accurate. More broadly, AI models are trained on general architecture that may not always provide useful results in the context of our proprietary codebase.

Question ASSET MANAGEMENT 4. What is the main barrier to developing AI in your use case (e.g. lack of skills and resources, readiness of the technology, high regulatory costs for compliance with the relevant frameworks, etc.)?

We do not have any significant barriers to benefiting from AI as part of our asset management processes. Use of AI-based tools is, and will continue to be, limited to methods that do not compromise information security. Nevertheless, our use and adoption of AI-based tools is progressing over time and will continue to do so.

Question ASSET MANAGEMENT 5. Does AI reduce or rather increase bias and discrimination in your use case?

Please explain and give examples when possible.

Asset management decisions are financial in nature and based on economic or statistical data, market conditions, and investable universe. Decisions made benefit all investors in the same strategy or portfolio equally.

In the context of asset management, statistical bias is a more pronounced concern than interpersonal bias or discrimination. It is possible that the use of AI may increase statistical bias in some cases, and lead to overestimating the probability of events that have already occurred.

This has long been a concern in investment research processes, which worry about overfitting. In that sense this is not a new concern, but rather a natural progression as more advanced systematic research and development processes come to light.

Question ASSET MANAGEMENT 6. Has general purpose AI opened new possibilities or risks in your use case?

- Yes

Like other technologies before it, AI-based tools allow for more efficiency and computational ease. Development cycles can be faster, debugging resolved more quickly. Research papers and other bodies of public information can be summarized quickly to help identify new areas of interest. While all of these things were possible previously, the time and costs to execute the R&D process is improved with the assistance of AI.

AI models can also potentially identify relationships in data that do not have clear economic intuition or human interpretability. This can potentially add nuance to asset management decisions, above and beyond what is solely possible from a fully human-directed process.

AI-based tools pose additional new risks of information leakage that must be carefully managed, however. Understanding the type of AI tool being used, and whether input data is provided back to the model as training data, is crucial to managing the risks the tool may pose to our business (and informs our decision on whether or not to use it).

Question ASSET MANAGEMENT 7. On whom do you rely for the development of your AI solutions?

- Partial collaboration with external providers

Please explain and give examples when possible.

We make internal use of general purpose LLMs provided by external parties. We also make use of open source libraries that implement machine learning techniques (*eg.* scikit-learn, PyTorch). These tools are integral to our research and development process, and assist us as we develop in-house software and applications, which may incorporate AI-enabled technologies.

All end decision making and adjustments to our internal infrastructure remain subject to human review, approval, and implementation.

We do not provide separate AI-based asset management services directly.

Question ASSET MANAGEMENT 8. When delegating functions to third parties, do you check the extent to which the provisions of services will entail the use of AI?

- Yes

Please explain and give examples when possible.

All our vendors go through a rigorous vetting process as part of their onboarding. Vendors are also subject to periodic security reviews as part of our regular course of business.

We inquire about how and where vendors may use AI in their business as part of their evaluation. This helps us understand how to best benefit from their services, as well as where there may be heightened risks.

Most importantly, we look to ensure that we maintain reasonable safeguards around our proprietary information and intellectual property, and adhere to business practices that afford us the benefit of third party's use of AI without exposing ourselves to undue risks.

We generally do not delegate asset management functions to third parties, although we may work with external parties and vendors to support parts of our asset management process. For example, we work with external developer teams (*eg.* LimeChain, Talos) to provide software as a service, or direct development support. These external teams may use AI-based tools like GitHub Copilot to aid in their development work.

In cases where we delegate function or service provision to third parties, we place particular scrutiny on how and where data is stored. We have a strong preference for all software and data to be stored in CBAM-managed locations with security oversight. This helps ensure that sensitive information and ideas are not at risk of compromise, either implicitly or explicitly.

3.1. Scope and AI definition

Question 34. Which of the following use cases that could fall into the categorisation of high-risk are potentially relevant to your activity?

- None.

Question 35. Please explain the overall business and/or risk management process in which the high-risk use case would be integrated and what function exactly the AI would carry out.

N/A

Question 36. Are there any related functions AI would carry out which you would suggest distinguishing from the intended purpose of the high-risk AI systems in particular to the use cases identified in question 34?

No.

Question 37. Please explain why these functions would/should in your view not be covered by the high-risk use cases set out in the AI act either because they would not be covered by the definition of the use case or by relying on one of the conditions under article 6(3) of the AI Act and explaining your assessment accordingly that the AI system would not pose a significant risk of harm if:

N/A

Question 38. At this stage, do you have examples of specific AI applications/use cases you believe may fall under any of the conditions from article 6(3) listed above? Please describe the use case(s) in cause and the conditions you believe they may fall under.

As a financial institution, we receive numerous requests for information from regulators or law enforcement agencies to assist with their investigations. We utilize machine learning to perform the narrow, procedural task of collating documents, which will be reviewed by humans prior to dissemination of the information. The ML tool's sole function is to synthesize the request details and gather documents. No assessments, decisions, or external responses are generated or submitted by the ML tool.

Question 39. Based on the definition of the AI system, as explained above (and in article 3(1) and accompanying recitals), do you find it clear if your system would fall within the scope of the AI Act?

- Yes

3.2. AI Act requirements

Question 40. Bearing in mind there will be harmonised standards for the requirements for high-risk AI (Mandates sent to CEN-CENELEC can be monitored here), would you consider helpful further guidance tailored to the financial services sector on specific AI Act requirements, in particular regarding the two high-risk AI use cases?

- Yes. If yes, on which specific provisions or requirements and on what aspects concretely?

Additional guidance would be helpful if it is tailored to specific use cases, rather than providing overarching principles. Because the EU AI Act (and similar AI regulations) are use-case based, implementing high level guidance would be unlikely to grant further clarity. Instead, specific clarifications—such as when "fine tuning" or "modifications" result in the creation of a new GPAIM—would be helpful. Similarly, understanding when a deployer crosses the threshold to be a downstream provider would be beneficial and would support companies' compliance efforts.

3.3. Financial legislation requirements

Question 41. Future AI high-risk use cases would also need to comply with existing requirements from the financial legislation. Would you consider helpful further guidance meant to clarify the supervisory expectations for these use cases?

- If yes, please explain your choice and indicate if the guidance should be high-level and principles based or tailored to specific use cases.
- No, the supervisory expectations are clear.

N/A

Question 42. There are other use cases in relation to the use of AI by the financial services sector which are not considered of high-risk by the AI Act, but which need to comply with the existing requirements from the financial legislation. Would you consider helpful further guidance meant to clarify the supervisory expectations for these use cases?

- If yes, please explain your response, and indicate if the guidance should be high-level and principles based or tailored to specific use cases.

As we discuss in our response to Question 40, guidance is best focused on specific use cases rather than high level principles. For example, guidance specifically granting scope to use AI and ML tools for fraud prevention and risk management would ensure that

companies can freely use these tools to benefit both users and compliance. Similarly, requirements for impact assessments, which are already being implemented under the EU AI Act, should recognize the benefits of these tools to consumers and not impose burdensome harm-mitigation requirements that risk undermining those benefits.

- No, the supervisory expectations are clear.

Question 43. Are you aware of any provisions from the financial acquis that could impede the development of AI applications (e.g. provisions that prohibit the use of risk management models which are not fully explainable or the use of fully automated services for the interaction with consumers)?

- If yes, please indicate the acquis/ provision in cause.

While we are not best positioned to speak to particular provisions, we note that AI tools and cloud adoption go hand in hand. Therefore, removing regulatory and supervisory barriers to cloud adoption would help facilitate responsible AI adoption.