



April 24, 2024

Mr. Christopher Kirkpatrick  
Secretary of the Commission  
U.S. Commodity Futures Trading Commission  
Three Lafayette Centre  
1155 21st Street, N.W.  
Washington, D.C. 20581

Re: Request for Comment on the Use of Artificial Intelligence in CFTC-Regulated Markets

Dear Mr. Kirkpatrick:

Coinbase Derivatives, LLC (“CDE”), Coinbase Financial Markets, Inc. (“CFM”), and Coinbase, Inc. (“CBI,” and collectively with CDE and CFM, “Coinbase”),<sup>1</sup> welcome the opportunity to respond to the U.S. Commodity Futures Trading Commission’s (“CFTC”) request for comment on “The Use of Artificial Intelligence in CFTC-Regulated Markets” (the “RFC”).<sup>2</sup> Coinbase fully supports effective regulation developed with the input and coordination of industry members and appreciates the thoughtful approach taken by the CFTC to better understand the implications of artificial intelligence (“AI”) systems on the CFTC-regulated marketplace. With entities registered with the CFTC as a designated contract market (“DCM”) as well as a futures commission merchant (“FCM”), Coinbase has endeavored to provide U.S. investors with access to derivatives markets in the United States in a regulated way that also helps keep the United States at the center of digital innovation. As CFTC-regulated markets evolve, CDE and CFM intend to continue supporting efficient and safe market operations while promoting product innovation and enhancing user experience.

Coinbase applauds the release of the RFC because it demonstrates concretely that the CFTC appreciates the impact that machine learning technology (“ML”)<sup>3</sup> will have—and is

---

<sup>1</sup> CDE operates a DCM while CFM is registered as an FCM with the CFTC and the National Futures Association (“NFA”).

<sup>2</sup> *Request for Comment on the Use of Artificial Intelligence in CFTC-Regulated Markets*, CFTC (Jan. 25, 2024) [https://www.cftc.gov/media/10156/AI\\_RFC\\_012524/download](https://www.cftc.gov/media/10156/AI_RFC_012524/download).

<sup>3</sup> 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.

already having—on the markets that it regulates. Indeed, Coinbase believes that ML methods will be a mission enabler for the CFTC, particularly as markets become digitally native while continuing a trend towards higher velocity and greater data generation.

Today, most of our experience with the ML methods underlying generative AI products and processes is associated with supervision of activities in CBI, which operates a digital asset spot exchange and holds 45 state money transmission licenses and a BitLicense from the New York Department of Financial Services. While CBI is not a CFTC registrant, the learnings are nonetheless relevant to the CFTC’s anti-fraud and anti-manipulation authority over spot commodity markets, and they are also relevant to supervisory practices that can be adopted in derivatives markets.

Greater adoption of AI-enabled technologies like those that we and other market participants are adopting will not only enable registered entities under the CFTC’s purview to better meet their own regulatory obligations but will also help the CFTC fulfill its own mission. 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, CFTC markets will enjoy higher integrity and offer greater safety for investors.

To be sure, the use of AI systems and tools is already present in CFTC-regulated markets. Systems using AI, as defined by the executive order issued by President Biden (the “Executive Order”),<sup>4</sup> are not yet widely deployed by Coinbase for product development and production, risk management, or other corporate functions. ML, which the Executive Order and RFC definition of AI captures,<sup>5</sup> is used by Coinbase in confined and carefully governed ways to improve and enhance processes and procedures implemented to perform certain functions for the company, with promising results.

The most prevalent use of ML systems at Coinbase is for trade surveillance programs on CDE and CBI platforms. Coinbase entities are also either using ML models or are exploring the use of AI systems to improve the overall customer experience. Finally, AI and ML models hold great promise in improving the customer onboarding process for the purpose of verifying customer identification and analyzing other data related to customers and their behavior in order to detect and prevent fraud. This comment letter discusses each of these areas below and, in doing so, addresses other questions posed by the RFC.

---

<sup>4</sup> Exec. Order No. 14110 (Oct. 30, 2023), 88 Fed. Reg. 75191 (Nov. 1, 2023) <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>.

<sup>5</sup> See *id.* at 75193, 75195 (defining “machine learning”); RFC at 1 (noting that the CFTC is “monitor[ing] the adoption of AI, including machine learning” in CFTC-regulated markets).

In summary, Coinbase wishes to convey the following key points:

- As more financial instruments are brought to market and CFTC markets continue to proliferate and evolve, transactions involving these products will continue to create enormous amounts of data. As a result, companies with regulatory compliance obligations will need to leverage ML and, eventually, AI tools in a responsible way to meet those obligations. Importantly, the immutability and transparency of public blockchain data will enhance these tools even further.
- As the CFTC is already aware, this transition will require that controls and governance addressing the use of these tools will be increasingly important to ensure such automated systems are analyzing and processing data appropriately and according to those regulatory obligations. The CFTC's recently proposed guidance on third-party relationship programs is a useful and adequate initial step in addressing the regulatory considerations presented by AI. We caution, however, that this is a nascent and rapidly developing technology whose promise and risks are best addressed in the CFTC's usual principles-based manner. A more prescriptive approach would most likely fail to keep pace with future developments in AI technology.

### **Discussion**

As a general matter, 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. Derivatives contracts are increasingly traded using programmatic methods, which leads to even more production of data as the number of trading messages increases in volume and velocity.<sup>6</sup> As a result, companies that offer such products will need to continue iterating their approach to risk management and compliance on their platforms in order to keep pace with this evolution. Leveraging automated tools for this purpose—including ML and, eventually, AI programs—will not be a luxury but, rather, a necessity in order for CFTC-registered entities to continue meeting their regulatory responsibilities.<sup>7</sup> Today's risk-management tools will not be adequate for tomorrow's markets.

---

<sup>6</sup> See John Coughlan & Alexei G. Orlov, *High-Frequency Trading and Market Quality: Evidence from Account-Level Futures Data* 1, 9 (July 29, 2022) [https://www.cftc.gov/sites/default/files/2022-08/HFT\\_and\\_market\\_quality\\_ada.pdf](https://www.cftc.gov/sites/default/files/2022-08/HFT_and_market_quality_ada.pdf).

<sup>7</sup> Speech by Mark Carney, Governor of the Bank of England, *Enable, Empower, Ensure: A New Finance for the New Economy* (June 20, 2019) <https://www.bankofengland.co.uk/-/media/boe/files/speech/2019/enable-empower-ensure-a-new-finance-for-the-new-economy-speech-by-mark-carney>. See also Jo Ann Barefoot, *The case for placing AI at the heart of digitally robust financial regulation*, Brookings (May 24, 2022) <https://www.brookings.edu/articles/the-case-for-placing-ai-at-the-heart-of-digitally-robust-financial-regulation/>.

## 1. Systems Using ML Models Can Assist with Combating Fraud, Bank Secrecy Act Compliance, and Trade Surveillance

CFM, which offers derivatives products to U.S. investors, is subject to the requirements of the Bank Secrecy Act (“BSA”) and must establish and maintain effective anti-money laundering (“AML”) programs to detect and prevent money laundering and other illicit activities.<sup>8</sup> As part of this program, the CFTC also requires that FCMs have robust customer identification procedures (“CIPs”), transaction-monitoring systems, and reporting mechanisms in place to identify and report suspicious activities.<sup>9</sup>

CBI is subject to similar obligations as a money transmitter and, in its continued efforts to enhance its BSA program, has leveraged ML systems where appropriate to improve its capabilities regarding BSA compliance and fraud prevention more generally. This letter describes below some of the important areas of this program where ML systems are used. In the future, similar ML systems could also facilitate compliance with CFTC registrants’ BSA obligations.

### a. Compliance During Onboarding and Customer Identity Verification

Verifying customer identity is a fundamental precursor to setting up a CBI account. CBI has policies and procedures that frame the company’s customer identification program, which enable the company to confirm the identity (“ID”) of potential customers seeking to onboard the CBI platform. These policies and procedures also govern the use of any third-party vendors that assist with the customer identification function.

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 be given 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 the submitted documentation and its association with the onboarding customer through a variety of different methods. It is during this onboarding stage where oftentimes 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.

---

<sup>8</sup> 17 U.S.C. § 42.2.

<sup>9</sup> 31 C.F.R. § 1026.220. Treasury, FinCEN, and the CFTC jointly issued 31 C.F.R. § 1026.220, which requires FCMs and introducing brokers (“IBs”) to have customer identification programs for identifying and verifying the true identity of customers. An FCM or IB’s written policies and procedures must enable it to form a reasonable belief that it knows the true identity of each customer.

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 CIP and the broader BSA program, all other considerations remaining equal.

b. Combating Fraudulent Behavior Post-Onboarding

Once an onboarding customer's identification is verified, there remain other risks related to fraud potentially presented during and after the onboarding process.<sup>10</sup> CBI has observed that certain data on the CBI platform serve as indicia of those risks, which include second-party and third-party fraud.<sup>11</sup> For example, those engaged in fraudulent conduct may sometimes change their name to similar ones, or to an alias, to avoid detection and will 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.

---

<sup>10</sup> See 7 U.S.C. § 6b. See also 31 C.F.R. § 1026.220.

<sup>11</sup> 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."

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 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.

c. Public Blockchains Can Improve the Utility of ML Models

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.<sup>12</sup> 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.<sup>13</sup>

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.<sup>14</sup>

---

<sup>12</sup> 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”).

<sup>13</sup> 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>.

<sup>14</sup> 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

#### d. Trade Surveillance

As market operators, CBI and CDE have implemented robust trade surveillance programs to detect potentially manipulative conduct on their platforms. CDE, Coinbase's DCM, is required to ensure fair and orderly trading on its platform through compliance with the CFTC's core principles, which include Core Principle 4, the "Prevention of Market Disruption."<sup>15</sup> Collectively, the core principles are designed to detect and deter market manipulation, fraud, and other abuses within the trading market itself. More specifically, DCMs must implement a trade surveillance program to monitor trading activity and detect and investigate any such activity indicating manipulative or fraudulent conduct.<sup>16</sup>

Many DCMs process millions of messages per trading day, reflecting vast amounts of data being created on the exchanges and creating the need for automated programmatic tools to implement an effective trade surveillance program. The same is true for CBI. More and more, these tools can include ML models, which can be designed to detect manipulative trading activities, such as "spoofing" or "layering," through ingestion of trade message patterns indicative of these stratagems. Such an ML surveillance model 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 has begun 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

---

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.

<sup>15</sup> Section 5(d)(4) of the Commodity Exchange Act, 7 U.S.C. § 7(d)(4), entitled "Prevention of Market Disruption," requires that a DCM "shall have the capacity and responsibility to prevent manipulation, price distortion, and disruptions of the delivery or cash-settlement process through market surveillance, compliance, and enforcement practices and procedures, including . . . (A) methods for conducting real-time monitoring of trading; and (B) comprehensive and accurate trade reconstructions." See *also* 17 C.F.R. § 38.250.

<sup>16</sup> "A designated contract market must maintain an automated trade surveillance system capable of detecting and investigating potential trade practice violations. The automated system must load and process daily orders and trades no later than 24 hours after the completion of the trading day. In addition, the automated trade surveillance system must have the capability to detect and flag specific trade execution patterns and trade anomalies; compute, retain, and compare trading statistics; compute trade gains, losses, and futures-equivalent positions; reconstruct the sequence of market activity; perform market analyses; and support system users to perform in-depth analyses and ad hoc queries of trade-related data." 17 C.F.R. § 38.156.

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.

In the future, CDE intends to use ML 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.

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.

e. Books and Records

As a DCM and FCM, respectively, CDE and CFM have a range of books and records obligations related to customer activities on their platforms in addition to related internal policies and procedures. Much of the data referred to in the foregoing discussion of CDE's trade surveillance program is implicated by these obligations, thus requiring that the data be captured and retained pursuant to applicable CFTC regulations. As explained, ML models can ingest trade data as it is being created in real time, allowing CDE to detect suspicious activity and automate responses to such activity nearly instantaneously. In this respect, and responding to a specific question in the RFC, ML models indeed are being used to "proactively search for risks in records and recordings," even though in this context such models were not built specifically to assist with CFTC books and records requirements applicable to CDE.<sup>17</sup>

As the RFC suggests, record retention requirements enable post hoc review of suspicious activity as well. Coinbase has not used AI tools for post hoc reviews of books and records in order to search for "gaps" therein as a general matter, but rather has focused its use of ML models on the particular types of fraud prevention and trade surveillance alluded to above. This includes post hoc review of transactions. Additionally, and as a matter of course, CBI does regular code reviews of all of its risk management and compliance systems to ensure

---

<sup>17</sup> See RFC Question 2e.



these systems are working properly and, in so doing, reviews all relevant books and records necessary to facilitate this review.

## **2. Systems Using ML Models Can Improve the Customer Experience**

Using ML programs to improve the experience for users of consumer products is relatively commonplace. Most readers of this comment letter will relate to the experience of an email or texting platform suggesting words or sentences as they type a message. CBI similarly uses ML models to improve the user's customer experience. An important element of the customer experience for CBI users is the home or landing place reached after signing in where information feeds can be tailored to the customer's preferences. ML models can learn to curate and deliver automatically the types of information that data shows a given customer prefers.

Likewise, by consuming and processing data generated from customer activity on Coinbase platforms, ML models might learn to alert a particular customer to a price movement in a particular asset or make the customer aware of similar products or investment opportunities, subject to any relevant regulatory requirements.<sup>18</sup> Coinbase also has explored ways that ML tools could improve the level of customer service that users receive after they identify issues that need to be resolved. For example, ML tools could generate personalized responses to specific customer queries conveyed through a chat function or other communication channel on the Coinbase platform. Coinbase has also considered whether AI and ML models could help guide a customer through the onboarding process as well as navigating the various platforms.

## **3. Managing Third-Party Vendors That Leverage AI Solutions**

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.

In the future and if finalized, the CFTC's recent rule proposal on operational resilience, including the proposed guidance on the use of third-party vendors, would apply to Coinbase's FCM vendor relationships, whether AI is involved or not.<sup>19</sup> In the meanwhile, Coinbase has

---

<sup>18</sup> Such requirements include 17 C.F.R. Part 160; 17 C.F.R. § 38.1051(a)(2).

<sup>19</sup> See *Operational Resilience Framework for Futures Commission Merchants, Swap Dealers, and Major Swap Participants*, 89 Fed. Reg. 4706 (Jan. 24, 2024) <https://www.cftc.gov/sites/default/files/2024/01/2023-28745a.pdf>. The proposal includes a proposed Appendix A to Part 1, entitled "Guidance on Third-Party Relationship Programs." See also NFA Interpretive Notice 9070, *NFA Compliance Rules 2-9, 2-36 and 2-49: Information Systems Security Programs* (Sept. 30, 2019) <https://www.nfa.futures.org/rulebooksql/rules.aspx?RuleID=9070&Section=9>; NFA Interpretive Notice 9079, *NFA Compliance Rules 2-9 and 2-36: Members' Use of Third-Party Service Providers* (Sept. 30, 2021) <https://www.nfa.futures.org/rulebooksql/rules.aspx?Section=9&RuleID=9079>. Recent guidance from the federal banking agencies is also instructive. See OCC Bulletin 2023-17, *Third-*

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.

Finally, any customer of an FCM or DCM must do its own due diligence when onboarding to such a platform, and such diligence might include a review of the FCM or DCM's own vendors. In this context, an FCM or DCM should disclose the relevant features of their

---

*Party Relationships: Interagency Guidance on Risk Management* (June 6, 2023)  
<https://www.occ.gov/news-issuances/bulletins/2023/bulletin-2023-17.html>.

vendor products when asked, including any AI systems. At this stage of maturity for AI systems, however, it is unclear that there would be much value in otherwise requiring that an FCM or DCM proactively and specifically disclose any use of AI systems. Such a mandatory disclosure might introduce other risks that could reduce the effectiveness of AI tools while not necessarily improving regulatory outcomes.

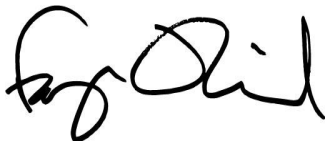
\* \* \*

## Conclusion

Coinbase uses ML models to improve outcomes and efficiencies in its various risk-management programs, including its BSA and trade surveillance programs. ML models also have helped Coinbase make strides in providing the best user experience possible for its customers. In some instances, Coinbase's vendors also use ML tools in their own products, requiring greater transparency by the vendor and deeper due diligence and monitoring by Coinbase, as the end user.

The CFTC's recent proposed guidance is a welcomed and adequate first step in addressing the risks posed by such vendor-delivered AI deployments—the proposal suggested useful recommendations for best practices to address those risks. More broadly, this RFC's attempt to better inform the CFTC is not only wise, but crucial—AI deployments by the CFTC's registered entities and registrants will be needed to meet the rising regulatory challenges posed by ever-more-data-generating, digitized markets. The CFTC's efforts to stay current on these developments and routinely assess whether its own regulatory programs are keeping pace with these developments will help ensure the agency is meeting the goals of its own mission.

Sincerely,



Faryar Shirzad  
Chief Policy Officer



Gregory Compa  
Senior Director, Head of Institutional  
Compliance