June 29, 2021

Board of Governors of the Federal Reserve System
Consumer Financial Protection Bureau
Federal Deposit Insurance Corporation
National Credit Union Administration
Office of the Comptroller of the Currency

Re: Theta Lake, Inc.'s Response to the Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning

To the Above-Listed Agencies:

Theta Lake, Inc. ("Theta Lake") submits this letter in response to the Board of Governors of the Federal Reserve System, Consumer Financial Protection Bureau, Federal Deposit Insurance Corporation, National Credit Union Administration, and Office of the Comptroller of the Currency's (collectively, the "Agencies") Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning (the "RFI").

This letter collectively responds to questions one (explainability), seven (cybersecurity), and ten (third-party risk), posed in the RFI. This letter will:

- discuss how Theta Lake uses Artificial Intelligence ("AI") in its security and compliance platform;
- suggest an approach to the Agencies' assessment of AI based on a set of high-risk categories; and
- discuss standard practices relating to cybersecurity, operations, and explainability that would facilitate strong AI development and aid in the evaluation of AI technologies.

Theta Lake recommends that the Agencies take a risk-based approach to any guidance relating to the development, procurement, and use of AI by financial institutions ("FIs"). Such an approach will allow FIs to assess AI technologies based on the concrete risks posed to the organization, requiring deeper engagement when the AI is used for a pre-defined high-risk activity. Defining high-risk use cases for AI such as underwriting, investment advice, or capital management, or the presence of certain high-risk AI inputs like gender or race, that necessitate robust assessment prior to use would facilitate more meaningful and efficient appraisal and deployment of AI technologies.

In addition, Theta Lake recommends several standard practices that organizations developing AI can take to evidence minimum administrative and security controls such as SOC 2, Type 2 audits, baseline explainabiltiy documentation, policies, and more.

It is our hope that these observations and recommendations can positively contribute to the ongoing dialogue about the use of AI in financial services, which ultimately impacts consumers, FIs, startups, agencies, and regulators.

Theta Lake's Use of AI

Of the many uses cases for AI, helping humans navigate the ever-growing volume of electronic communications is a powerful one. With the need to review communications from documents and emails to chats, texts, voice calls, and video meetings, human reviewers in compliance, security, and regulatory oversight functions have an uphill task in attempting to watch, read, and listen to those conversations. AI can assist in this manually intensive review process.

Theta Lake uses AI in the form of machine learning ("ML") and natural language processing ("NLP") to analyze the increasingly dynamic and diverse data from collaboration, chat, audio, and email applications that enable communications between FIs and their employees, customers, agencies, regulators, and advisors. Platforms like Zoom, Slack, Microsoft Teams, and Cisco Webex are the foundations of the dispersed workforces that have emerged during the COVID-19 pandemic and will remain the bedrock of permanent hybrid work environments and the office of the future.

Theta Lake's cloud-based platform ingests the spoken, shown, and shared elements of collaboration and chat communications through direct, API-based integrations. Once ingested, Theta Lake uses ML and NLP to analyze video, audio, and text content for regulatory, security, and privacy risks.

Theta Lake uses ML and NLP in conjunction with computer vision and optical character recognition to examine content displayed over screen shares, webcams, and native whiteboards. These techniques are also used to analyze images, gifs, and reactions that comprise modern chat conversations. Transcribed audio conversations, text from chat communications, and files of all types transferred during conversations flow through our ML and NLP analysis pipeline and are scanned for risk. The use of ML and NLP to examine these content types facilitates identification of risk across the dynamic communication components of modern collaboration and chat applications.

Following analysis of the content, identified risks are displayed in an intuitive review screen to allow compliance and security teams to engage directly with potential issues. It is important to stress that human review by a compliance or risk team is the last step in the supervisory process—Theta Lake does not make automated decisions about potential risks or rule violations. Theta Lake highlights items of potential interest in context for a human consideration.

Basically, the AI components of Theta Lake's platform facilitate an understanding of the context of a conversation and result in increased review efficiency by flagging portions of conversations that require further analysis by an individual.

AI is incorporated into Theta Lake's built-in risk detections that examine the video, audio, chat, and file transfers in collaboration and chat conversations for universally common concerns in communications. Theta Lake has developed roughly 70 risk detections, which are pre-trained and ready for customer use with customers able to provide feedback and training on these classifiers. Customers can also engage Theta Lake to create customized AI-based risk detections for issues relevant to their specific product offerings, business units, or security and compliance concerns.

By way of example, Theta Lake deploys AI to identify compliance issues related to customer complaints under FINRA Rule 4513 and CFPB Rules, self-promotional language under Regulation Z and FINRA Rule 2210, and derivatives transactions relevant to CFTC Regulation § 23.202. Theta Lake also uses AI to detect the presence of personally identifiable information like names, email addresses, birthdates, or account numbers displayed on screen, spoken during conversations, typed in a chat, or included in file transfers. From a security perspective, AI is used for risk detections that examine screen shares for malware URLs, the display of sensitive applications like Quickbooks, Gutso, or Zoho as well as financial logos, adult brands, hate speech, and other offensive or contentious content.

Theta Lake's application of AI to these risk domains provides FIs full transparency into interactions taking place on their collaboration, chat, and audio systems. The ability to identify potentially problematic conduct protects consumers, facilitates regulatory compliance, mitigates leakage of sensitive data, and enhances security practices.

A Risk-Based Approach to Evaluating the Use of AI in Financial Services

While Theta Lake deploys AI for risk detection in the context of communications platforms, the implementation of AI in financial services spans a diverse array of use cases, including making determinations about creditworthiness, capital adequacy, customer support, and recruiting. Given AI's expanding use, a key pillar of any future guidance regarding FIs' assessment practices for AI should be grounded in a risk-based approach.

The Agencies should provide basic guidance about the kinds of activities that may be considered high-risk and develop a principles-based approach that outlines considerations for the assessment AI technologies based on the presence of high-risk factors. A regulatory approach predicated on risk assessment will allow financial institutions to adopt AI technologies by applying an appropriate level of scrutiny derived from the measurable potential business, financial, and customer impacts of a given use case. A flexible, risk-based framework for AI assessment will facilitate more meaningful and efficient development and assessment of AI solutions overall.

All risks are not created equal and a multitude of overlapping risks could arise in a particular context. Some risks may be exclusively financial in nature while others may arise from regulatory, operational, reputational, market, or credit-based factors. Given the varied nature of risks and impacts, any guidance regarding the assessment protocols for AI should be rooted in the magnitude of potential risk.

In our experience, a one size fits all approach for assessing AI results in a process-driven analysis that, in an attempt to answer a standard set of questions, fails to account for the proposed use of the underlying technology and meaningfully tailor the assessment based on practical risks. An approach guided by the purpose for which the AI is being deployed, informed by a set of high-risk categories, would facilitate easier identification of potential issues in AI technologies as well as a more efficient analysis process, allocating review resources based on risk severity.

Existing guidance like the Federal Reserve's SR 11-7: Guidance on Model Risk Management is already being applied to AI technologies. And although SR 11-7 refers to the notion of materiality as the basis for model evaluation, high risk and materiality have important distinctions. While materiality considers the impact or importance of a given model generally, an approach informed by specific high-risk categories would provide an explicit roadmap for organizations developing AI. Approaches based on materiality may result in significant efforts expended to exhaustively scrutinize AI applications that may present drastically different risks to an organization.

We implore the Agencies to consider the adverse impact of a singular approach to AI evaluation processes.

Examples of High-Risk AI Categories and Attributes

A preliminary designation of the following activities as high-risk may be useful as a starting point given their potential impact to consumers and related FI processes and controls:

- underwriting;
- valuing exposures, instruments, and positions;
- managing and safeguarding client assets;
- determining capital and reserve adequacy;
- investment advice;
- human resources management, including hiring, promotion, or termination; and
- consumer lending, banking, or credit determinations.

Similarly, the use of the following data attributes in developing AI, or as the basis for AI-informed decisions, might be considered high-risk factors:

- race;
- color;
- religion;
- national origin;
- sex;
- familial status;
- age;
- ethnicity;
- disability;
- genetic information; or
- any other socioeconomic or demographic information.

For example, an AI technology used to assess the age, gender, income, and race of a loan applicant would require a more comprehensive review than an AI application used to analyze office energy efficiency.

These lists are not intended to be exhaustive and operate as a starting point to define the types of financial activities and data points that, if incorporated into an AI platform, would be designated as high risk and require a more comprehensive assessment prior to deployment.

The European Commission's Approach to High-Risk AI as a Blueprint

Using a risk-based approach to inform AI evaluation has found favor in global frameworks. For example, the European Commission's ("EC") proposal "Laying Down Harmonized Rules on Artificial Intelligence and Amending Certain Union Legislative Acts" 2021/0106 (COD) adopts a risk-based approach to evaluating AI. In the EC's proposed implementation:

> In order to introduce a proportionate and effective set of binding rules for AI systems, a clearly defined risk-based approach should be followed. That approach should tailor the type and content of such rules to the intensity and scope of the risks that AI systems can generate. (Paragraph 14, p. 21)

Moreover, the EC defines a set of high-risk activities that, by default, require a more rigorous analysis. (*See* Annex III – High-Risk Systems Referred to in Article 6(2)). High-risk systems include, but are not limited to, law enforcement, migration, asylum, and border control management. Financial services are encompassed in the set of high-risk systems as part of a category related to "[a]ccess to and enjoyment of essential private services and public services and benefits," which includes:

> AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems put into service by small scale providers for their own use; (Annex III, Section 5(c))

Although the scope of the EC's proposed rules are broader than financial services, the Agencies should consider a similar principles-based approach driven by high-risk thresholds when considering guidelines for the assessment of AI.

Best Practices for Organizations Developing AI

We believe there are a set of uniform best practices that organizations developing AI can implement, so that they can be held accountable for adhering to pre-defined standards, regardless of their risk profile. Protocols such as annual auditing and documentation standards should be implemented by any organization developing AI to mitigate cybersecurity and third-party risk. The following briefly highlights practices for consideration as mandatory controls for AI development.

Annual Security Audits

As a threshold matter, organizations developing AI must be able to demonstrate robust internal practices as they pertain to enabling appropriate technical and administrative security and privacy controls. Conducting routine annual audits such as the SOC 2, Type 2 or ISO 27001 meaningfully test controls that protect data. These audits include key cybersecurity components such as managing administrative access, vulnerability scans, penetration tests, incident response plans, vendor management, and the maintenance of written information security programs.

Annual audit processes like SOC 2 and ISO sit at the intersection of cybersecurity and third-party risk, providing repeatable, measurable protocols that demonstrate compliance with articulated controls, including those key to AI development processes. Since these audits collect a set of uniform controls and apply them to companies of all shapes and sizes, they offer a consistent metric with which to assess supply chain and third-party risk across organizations developing AI technologies.

Policies and Procedures

The creation and enforcement of policies and procedures such as Codes of Conduct, Acceptable Use Policies, and Information Security Polices indicate an organization's approach to transparency and accountability, which can be extended to management of AI development. These policies are typically required as essential controls under SOC 2 and ISO regimes. Since such policies broadly include requirements for data collection, use, and retention, they offer insight into an organization's approach to developing AI. Moreover, policies around machine learning operations and other AI-specific processes are increasingly common.

Including analysis of these documents as part of AI due diligence will provide FIs and the Agencies with important background information to assess program maturity and the effectiveness of an organization's controls related to AI-specific risks.

Baseline Explainability Documentation

The development and maintenance of easy to understand documentation around AI explainability can be of great benefit. Drafting a description of how an organization's AI has been developed that includes a plain language description of its functionality as well its practical use within a regulated organization is essential. Additional details regarding how AI is trained, which publicly available algorithms are employed, as well as high level details on testing and maintenance methodologies are also helpful.

Transparency and AI Performance Auditability as Part of System Design

In addition to the development of internal policies and security frameworks, considering transparency and auditability of AI as part of the software development lifecycle process is essential. While the purposes for which AI is used are varied, product features that include performance metrics and testing capabilities to oversee AI provide quantitative metrics that allow FIs to benchmark functionality and supervise AI-reliant processes.

For example, Theta Lake has developed a series of reports that provide metrics about its risk detections, which can be leveraged by customers for supervisory and testing purposes as well as to directly suggest re-training of classifiers and models. These reports provide tangible evidence of AI performance as well

as strengthen feedback loops between Theta Lake and its customers, resulting in closer working relationships and a stronger collective vision for platform development. Although, the engineering elements of AI platforms are often the focus of regulatory inquiry, we have found that supplemental features that encourage tighter business alignments are also extremely productive.

Conclusion

Given that standard compliance, risk, and security assessments leverage risk-based approaches, we recommend that the Agencies extend an equivalent evaluation paradigm to AI technologies. As AI is employed for a variety of internal and external use cases, a flexible and rigorous evaluation methodology informed by a set of high-risk use cases would provide clarity and predictability for FIs and companies developing AI-based technologies. An overly prescriptive regime would hamper innovation and negatively impact the efficiency and effectiveness of FIs, which would be required to redirect capital and resources to analyze AI regardless of its size, scale, or risk.

We would welcome the opportunity to discuss these issues with you further to offer the perspective of a security and compliance startup with experience building AI technologies for the financial services industry. Please do not hesitate to contact us with any questions.

Respectfully submitted,

Marc Gilman

Marc Gilman
General Counsel and VP of Compliance