February 2, 2023

Reporting, Disclosure, Data Strategy and AI Team
Prudential Regulation Authority
Threadneedle Street
London
EC2R 8AH

Submitted via email

Re: Theta Lake, Inc.'s Response to Discussion Paper DP5/22 – Artificial Intelligence and Machine Learning

To Whom It May Concern:

Theta Lake, Inc. ("Theta Lake") submits this response to the joint Financial Conduct Authority ("FCA") and Prudential Regulation Authority ("PRA") (collectively, the "Regulators") Discussion Paper 5/22 Artificial Intelligence and Machine Learning.

Below Theta Lake provides an overview of its Security and Compliance Suite and describes how artificial intelligence ("AI") is incorporated into the platform. Theta Lake also responds to Questions 1, 3, 7, and 17 posed in Discussion Paper DP5/22.

About Theta Lake

Theta Lake unlocks electronic communications systems for the modern financial services workplace with archive connectors, compliant archiving, and AI-based regulatory, security, and privacy detections across video, voice, chat, email, and document content.

With support across over 100 platforms including Zoom, Microsoft Teams, Slack, Bloomberg, Symphony, WhatsApp, and Webex by Cisco, Theta Lake enables compliance and security controls to support the dispersed workforces that emerged during the COVID-19 pandemic and that are the bedrock of hybrid work environments and the office of the future.

Theta Lake's ability to capture and analyze data from documents and emails to chats, texts, voice calls, and video meetings, allows firms to deploy their communications platforms of choice and facilitate AI-driven compliance and security oversight at scale. Theta Lake's communications platform integrations coupled with its sophisticated AI risk detections transforms the use and oversight of e-comms applications facilitating better and more efficient compliance.

Theta Lake's Use of AI

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 video, voice, chat, and email applications that enable unified communications between financial services firms and their employees, customers, agencies, regulators, and advisors.

Theta Lake's cloud-based platform ingests the spoken, shown, and shared elements of collaboration, voice, chat, and email communications through direct, secured application programming interface

("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 whiteboards. These techniques are also used to analyze images, gifs, and reactions that comprise modern chat and collaboration conversations. Transcribed audio conversations, text from chat communications, and files of all types transferred during interactions 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. The ability to examine these various data types enables more effective oversight of communications and allows firms to open up applications and feature sets like those described above offering employees an improved digital experience that reduces the use of non-approved communications channels.

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 review by a compliance or security professional using a "human-in-the-loop" as the key final step in the decision-making process.

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. The result is a workforce encouraged and excited to use firm-approved communications channels, reducing non-compliant off-channel conversations that introduce risks to customers, counterparties, and firms, and result in significant financial and remedial sanctions.

Theta Lake's AI-Enabled Risk Detections

Theta Lake has developed over 85+ AI-enabled risk detections, which are pre-trained and ready for customer use. As part of the human-in-the-loop process, users can provide feedback and training on the results of Theta Lake's risk detections to further refine them. 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, market abuse, and data protection such as detecting the presence of sensitive personal 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 HR or finance tools as well as the presence of financial logos, adult brands, hate speech, and other offensive or contentious content.

In the FCA context, Theta Lake uses AI to detect statements that would be considered misleading or inappropriate promotions or advice risks under COBS 4.2 and other relevant rules. Additionally, Theta Lake's AI can be used to detect if the Key Information Document required under the PRIIP regulation has been displayed, discussed, or transferred during a Zoom, Teams, Slack, or other collaboration communication.

Theta Lake's Compliance, Security, and Privacy Controls

Theta Lake's application of AI to these risk domains provides financial services firms with full transparency into interactions taking place across their unified communications, 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.

Specifically, Theta Lake's selective archiving, chat connectors, and AI-powered Suite enable firms to compliantly and securely unlock all the features in unified communication platforms which means that, in practice, there is less chance of employees deliberately avoiding monitored channels. As Theta Lake's fourth annual survey report found, 94% of firms are using up to six different communications applications, with 67% expecting those numbers to increase during the next 12 months—signaling a pressing need for modern compliance and security controls. Based on recent news reports, it is our understanding that the FCA is actively engaged in potential enforcement in this area, so innovations like Theta Lake's AI-enabled detections are an essential component for a future-proof compliance program.

Given the FCA's electronic communications recordkeeping and supervision mandates under several rules including SYSC 10A.1.6, MiFID II, and MAR, firms must contend with the oversight of exponentially increasing volumes of communications data. Moreover, the FCA's guidance in Market Watch #66 (2021) reinforced firms' recordkeeping and supervisory obligations as extending to remote work environments. Since supervision of e-comms is core to the FCA's overarching goals of market transparency and consumer protection, ongoing engagement and leadership in this area, including understanding of innovative AI-based compliance tools, is essential.

Theta Lake's Responses to Questions 1, 3, 7, and 17

Q1: Would a sectoral regulatory definition of AI, included in the supervisory authorities' rulebooks to underpin specific rules and regulatory requirements, help UK financial services firms adopt AI safely and responsibly? If so, what should the definition be?

A sectoral regulatory definition of AI is unlikely to help UK financial services firms adopt AI safely and responsibly. Given the potential impact of AI across industries, any financial services-specific definition is likely to be problematic as it would be difficult to remain technology neutral whilst also avoiding an unduly "tick box" approach to the assessment and considerations of a principles- and risk-based approach to AI. Further, an overly prescriptive regime based around a sectoral regulatory definition would be in danger of hampering innovation and negatively impact the efficiency and effectiveness of firms considering use case(s) for AI.

Q3: Which potential benefits and risks should supervisory authorities prioritise?

A principles- and risk-based approach to the adoption of AI in UK financial services is a more suitable approach. While Theta Lake deploys AI for risk detection in the context of unified 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 firms' assessment practices for AI should be grounded in a risk-based approach.

The Regulators should seek to 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 of AI technologies based on the presence of high-risk factors. A regulatory approach predicated on risk

assessment will allow firms 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.

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.

An approach guided by the purpose for which 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 vetting process, allocating assessment resources based on risk severity. Such an approach will also allow firms 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, capital management, or the presence of certain high-risk inputs like gender or race, that necessitate robust assessment prior to use would facilitate more meaningful and efficient appraisal and deployment of AI technologies.

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 examples are not intended to be exhaustive and should be seen 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.

Q7: What metrics are most relevant when assessing the benefits and risks of AI in financial services, including as part of an approach that focuses on outcomes?

Objective, data-driven metrics should be the Regulators' focus when assessing the benefits and risks of AI in financial services. Evaluating AI based on consistent and predictable outcomes will promote its use in a suitable and sustainable manner.

At Theta Lake, we've taken several steps to ensure that customers have transparency into the performance metrics of our AI-based risk detections and mechanisms for providing feedback for potentially anomalous behavior. Our classifier audit report provides metrics around positive hit rates for a given risk detection so, for example, customers can query the number of Zoom video meetings triggering the detection for misleading or inappropriate promotions in the video, voice, chat, and file transfer communications across a specific time range or group of users.

In addition to metrics around the performance of AI, Theta Lake also focuses on the performance of its models in the long term to minimize data and concept drift. Models are monitored post-deployment for unexpected changes and incorporate feedback from customers around false positives and false negatives. Models are refined and retrained routinely over time, incorporating new data sources in training data sets to account for emerging patterns.

When considering relevant metrics, the Regulators should consider both the benchmarks available to the users of an AI-based platform as well as the internal, operational practices of those developing the AI solution. Regardless of the implementation, AI cannot be a "set it and forget it" exercise.

While the metrics above may not apply in every AI use case, the broad concepts of ongoing oversight and supervision of IT systems are generally covered in the FCA's PS21/3, the Operational Resilience Requirements included in SYSC 15A.2, the May 2021 Implementing Technology Change Review, and Principles for Businesses (PRIN 2.1), Principle 3 mandating firms to take reasonable care to organize and control its affairs responsibly and effectively, with adequate risk management systems. These principles are also embodied in the PRA's PS6/21. These FCA and PRA requirements demonstrate how existing rulesets and reviews can be adapted to address potential risks in emerging technologies like AI.

Q17: Which existing industry standards (if any) are useful when developing, deploying, and/or using AI? Could any particular standards support the safe and responsible adoption of AI in UK financial services?

The issues of cyber security and data privacy must be considered with regard to the burgeoning use of AI. Indeed, organizations developing AI should 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, the software development lifecycle process, 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 critical 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 developing AI technologies.

Moreover, privacy by design standards that facilitate automated, scalable protection of sensitive personal information are critical. Theta Lake leverages AI to identify sensitive personal data like birthdates, national ID numbers, email addresses, and account numbers in visual, voice, and text content for the purpose of redacting, remediating, or removing them from conversations to minimize exposure and reduce the risk of misuse.

For example, Theta Lake's AI can identify a Webex screen share that includes a Microsoft Excel spreadsheet containing a list of National ID numbers and redact those numbers so that they are not visible to reviewers. In another use case, Theta Lake's AI can identify a Slack conversation containing email addresses and birthdates and remove that content from the chat itself to prevent its proliferation among participants, while retaining the original data in the background should it be required for regulatory or litigation purposes.

In terms of external references, Theta Lake finds the Financial Markets and Standards Board spotlight review of Monitoring FICC Marketing and the Impact of Machine Learning as a helpful contextual document.

Theta Lake would welcome the opportunity to discuss these issues further to offer additional perspectives from a fintech with experience building AI technologies for the financial services industry. Please do not hesitate to contact us with any questions.

Respectfully submitted,

/s/ Marc Gilman
Marc Gilman
General Counsel and VP of Compliance

/s/ Susannah Hammond
Susannah Hammond
Senior Regulatory Intelligence Expert

/s/ Stacey English
Stacey English
Director of Market Intelligence

/s/ Sharon Hüffner
Sharon Hüffner
Chief Scientist