June 25, 2021

Via electronic submission

Chief Counsel's Office
Attention: Comment Processing
Office of the Comptroller of the Currency
400 7th Street, SW
Suite 3E-218
Washington, DC 20219

Ann Misback
Secretary Board of Governors of the Federal Reserve System
20th Street and Constitution Avenue NW
Washington, DC 20551

Comment Intake
Bureau of Consumer Financial Protection
1700 G Street, NW
Washington, DC 20552

James P. Sheesley
Assistant Executive Secretary
Attention: Comments-RIN 3064- ZA24
Federal Deposit Insurance Corporation
550 17th Street NW
Washington, DC 20429

Melane Conyers-Ausbrooks
Secretary of the Board
National Credit Union Administration
1775 Duke Street
Alexandria, VA 22314-3428

Re: Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, including Machine Learning

Dear Sir or Madam:

Thank you for the opportunity to comment on the use of artificial intelligence (AI) tools at financial institutions, including areas like risk management and underwriting. As a company that provides integrated risk management and compliance software to a rapidly expanding customer base of over 3,600 financial institutions, Ncontracts has witnessed firsthand the burgeoning impact of AI on our customers and the

potential it holds for risk management, compliance, and fair lending, among other areas. AI presents both risks and opportunities, especially for community financial institutions, and it's important to always be looking ahead to ensure the financial system is prepared, protected, and positioned to benefit from advancements in technology.

Question 9: Do community institutions face particular challenges in developing, adopting, and using AI? If so, please provide detail about such challenges. What practices are employed to address those impediments or challenges?

Community FIs developing, adopting, and using AI face an uphill battle for three key reasons: regulatory challenges; lack of resources; and data management difficulties.

Regulatory challenges. Artificial intelligence is a new and emerging field where laws and regulation haven't yet caught up. As a result, FIs that want to implement AI must operate in a hazy area. There simply isn't clear guidance on which types of AI activities are permissible or how examiners will evaluate those activities. Meanwhile, examiners are trained to focus on the basics of banking during an exam. Very few have more than a cursory understanding of how AI works, and with no regulatory guidance for measuring or approving AI technologies, examiners are left scrambling when trying to assess an FI's use of AI—especially at smaller FIs.

That leaves FIs at a disadvantage when they try to justify why an AI model is appropriate for their FI or when demonstrating that their AI has been vetted and is not discriminatory. FIs must try to explain advanced-level AI to examiners who, while experts in FI operations, may not even have rudimentary AI knowledge. Understandably, examiners are uncomfortable giving the greenlight to a technology they don't fully understand, but that means a FI could potentially invest significant time and money into AI technologies only to walk them back if an examiner has trouble making sense of it. No amount of due diligence and transparency can overcome that obstacle. Examiners don't want to make a mistake, so they err on the side of caution, limiting innovation. Examiners are trained to review explicit formulaic methodologies. AI, while deterministic, lacks explain-ability and does not readily lend itself to direct examination. This will challenge existing models of examination.

Successfully overcoming this challenge requires work on the part of the regulatory agencies. The agencies and their examiners must expand their AI knowledge and training and provide more specific guidance explaining what examiners are looking for when assessing AI initiatives. They must concentrate on outcomes rather than the algorithms used to derive that outcome.

Lack of knowledge and resources. AI is a constantly evolving field. Many community institutions aren't prepared or comfortable keeping pace. They don't have the multi-million-dollar budget needed to hire developers to create AI or the software components needed to connect disparate systems. They often don't have the in-house expertise to even understand what types of third-party AI solutions are applicable to the FI's operations, how a particular type of AI works, or what controls are necessary to mitigate the risk.

Community FIs also have difficulty managing oversight of AI. Third-party AIs are often a black box. The FI doesn't entirely know how decisions are made because the third-party vendors that own the technology may not want to share proprietary data on how the AI draws its conclusions. For example, an AI could make credit decisions that result in discriminatory practices, and the FI wouldn't necessarily be aware of it or understand why it's happening. That poses a very real compliance risk, especially when using less mature products.

AI is also expensive, especially when it isn't scaled. Large FIs with the funds and expertise to implement AI have a competitive advantage over smaller FIs, and this advantage continues to grow as large FIs adopt increasingly advanced AI.

Disorganized data. AI runs on data, but the data at community financial institutions—and even large institutions—cannot easily be leveraged for this purpose. While there is no shortage of data, that data isn't in the right form or the right place. Data is often segregated, stored in different applications and systems that can't easily speak to each other. Before the data can be plugged into an algorithm to train AI, it requires cleaning, tagging, and filtering to be useful. It can be an extremely manual and time-consuming process that is just as tricky today as it was 50 years ago. It's a problem that neither large or small FIs have solved yet, and it's a significant obstacle to AI adoption.

There are also concerns about data security. The huge cache of data needed to train AIs is a tremendously appealing target for cybercriminals. Unauthorized access to this data, whether through a third-party AI provider or within the FI itself is a huge risk.

Finally, walled garden systems such as the FI's core system make data accessibility a challenge. These systems are purposefully opaque with respect to the data they contain, and vendors are reticent to give up this control.

Question 10: Please describe any particular challenges or impediments financial institutions face in using AI developed or provided by third parties and a description of how financial institutions manage the associated risks. Please provide detail on any challenges or impediments. How do those challenges or impediments vary by financial institution size and complexity?

When using AI developed or provided by third parties, the biggest challenges facing FIs are data and cost.

AI vendors often provide pre-built models created by pooling together a FI's data with data from other FIs. This is often necessary because a single community FI may not have enough data on its own to build a robust model—but it also means that the FI is relying on other FIs' data. If that data has problems or reflects unknown biases, those issues will influence the model.

Even when a third-party vendor uses the FI's own data to produce a model, the FI typically doesn't own the model. Vendors either do not want to explain their proprietary algorithms or how the data is crunched, or they cannot explain it due to the inherent nature of deep learning algorithms. (Small FIs are particularly at a disadvantage when seeking insights since they have little leverage.) This makes it hard for FIs to understand the model or demonstrate how it works to regulators. It also creates a business continuity issue. If the AI vendor is unable to perform, the FI will not have access to the model it has been using.

Consider the Bank Secrecy Act (BSA) requirement that an FI provide a risk rating for customers based on suspicious activity and currency transactions. When there is no access to a vendor's model, an FI will have a hard time ensuring the correctness of the model and won't be able to easily defend itself when an examiner challenges the model or the appropriateness of the FI's due diligence. Further, while most FIs spend tremendous amounts of resources on anti-money laundering (AML) because there have been huge fines associated with this space, but most community based institutions may not be able or willing to spend on AI in areas that considered profit centers, such as compliance or risk management. Costs are high and there are only so many resources to go around.

Third-party due diligence is also a challenge. Even if a vendor is willing to show how its model works, many FIs don't have staff with the expertise to understand how the models were created or identify any potential biases. Experienced data scientists are expensive and in short supply, making them out of reach for most FIs.

FIs also face data security risk when using a third-party AI provider. The computers needed to crunch the huge amounts of data required are very large, and most community FIs need to outsource this function. That means an FI must allow their data to leave the network.

For example, if an FI shares its complaint management data, that may involve allowing a vendor's AI to read every email that comes into the FI since a complaint can come from anywhere. This requires FIs to open up data to third-party vendors in new ways, potentially exposing even more sensitive data to increased risk. As a result, every AI vendor becomes a critical vendor. This risk is deepened if a third-party vendor outsources activities and shares the FI's data with a fourth-party vendor. FIs need to know their data exposure footprint and ensure protection in every location.

There are also questions about how to share the data. Will legacy systems be able to share data or will FIs need to build specific API integrations, which is an expensive investment? FIs will also have to decide if they are comfortable with letting AI partners integrate with their systems and cores (which is a significant data governance issue) or whether they prefer to push data so they can control which data leaves.

Question 13: To what extent do model risk management principles and practices aid or inhibit evaluations of AI-based credit determination approaches for compliance with fair lending laws?

Models are useful tools for evaluating the performance of AI-based credit decisions and the accuracy of assumptions, but there are also significant challenges. There is a real risk of bias when using AI-based credit determination.

Models are based on hypothetical data, and that data may contain biases that influence the results of AI. Consider the example of AI used in prison sentencing. When sentencing a prisoner, a judge is supposed to make decisions based on the data in front of them. Factors like their appearance, accent, or skin color should not be considered—yet it often is.

If an AI was built using real-life historical sentencing data, the AI would follow the example of the judges. It would be more likely to deliver harsher sentences to prisoners based on factors that correlate with things like appearance, accent, or skin color that should be discounted. For example, if it found that those with Latino last names were historically more likely to have longer sentences, it would give longer sentences to prisoners with Latino last names. It would discriminate just as much as a human judge.

Compare that to an AI that was built using historical and psychological data on recidivism. It was not given data on race, ethnicity, or other discriminatory factors. The machine was far more likely to give shorter sentences to people of color than judges with 20 to 30 years of experience.

Consider the same exercise using an FI's lending data. Credit decisions often have a certain amount of discretion built in. It's there to allow FIs to be flexible and creative with their lending, but it can also lead to discrimination as factors like appearance, accent, or skin color may inadvertently influence the lender's

decision. If you feed that data to a machine, you're perpetuating any built-in bias. A machine will find correlation between less favorable rates or terms or approvals and specific prohibited basis and use it to inform its decisions—unless you find a way to tell the machine how to avoid making decisions based on a prohibited basis. There need to be guardrails to steer the AI away from discriminatory and other unwanted results.

One way to create those guardrails is using AI with structured learning (i.e. learning that comes with instructions to prevent specific outcomes) versus unstructured learning (when the AI performs analysis and reaches conclusions based only on the data). Unfortunately, it can be hard to weed out unwanted results. The data used to train AI may contain unintentional biases. If that machine is used to train another machine and they feed the results back and forth, it will continue to generate those results in an unending feedback loop. Leveraging structured learning in modeling also requires a high-level of AI expertise that's typically not available to community financial institutions.

Another challenge is that the huge amounts of data make it difficult to identify which factors cause an AI to deny a loan application. Traditional lending models may base decisions on a defined list of distinct factors. When an application is denied, this model makes it easy to comply with regulation requiring that a denied applicant must receive an explanation that includes at least two important factors. An institution can easily show how the decision was made, outlining how the applicant fell short on two specific factors out of seven factors. AI models, especially those in a black box, make it much harder to explain denials to consumers. An AI model might determine that the applicant's coefficient factor of 3.6453 puts the applicant just outside of the mean in its 18 factors. That explanation will mean nothing to the applicant and will not be accepted by regulators either.

Fair lending is a high priority for the regulatory agencies, but its enforcement may be hampered if regulators are unable to decipher whether AI models and their underwriting decisions comply. Big tech has huge amounts of knowledge and much deeper pockets, giving it capabilities the government doesn't. How will regulators regulate AI models that provide essential services to FIs if they don't have a thorough understanding of how they operate?

Question 16: To the extent not already discussed, please identify any additional uses of AI by financial institutions and any risk management challenges or other factors that may impede adoption and use of AI.

AI has broad applications beyond underwriting. They include:

Fraud detection. One of the primary functions of AI is detecting and preventing fraud. Some larger financial institutions are doing this well, but many aren't. For example, some FIs block all transactions conducted abroad, even legitimate ones. This is a tool that must be improved and implemented at financial institutions, especially small ones, to protect consumers and reduce fraud risk. Additionally, model auto generation can help FIs identify outlier transactions that an FI might not have identified on its own.

Real-time risk trends. The FDIC collects call report data and by the time it's delivered to FIs, it's outdated. The 80-page report must then be read and assessed to determine how it impacts the FI's risk assessments or risk appetite. By the time the FI has time to leverage this risk management data, a new report is out. This creates a cycle where FIs are reactive instead of proactive.

It would be immensely useful if call reports and other data collected by regulators could be fed into AI to produce real-time key risk indicators (KRIs) that let FIs know what to look out for based on which products, services, or business area the FI operates in. It would also be helpful to provide anonymized data about findings. This would require a partnership between the agencies, FIs, and the third-party vendors that provide the risk modeling. This data would feed into models, making them more insightful.

Complaint management. AI is also a valuable tool for compliant management, especially if it can be leveraged to identify complaints and determine the potential violation so that it can be routed to the right person or area promptly. If a FI can quickly collect, manage, and resolve complaints from various channels earlier, both the consumer and FIs will benefit. Consumers' problems will be resolved both individually and universally, allowing a FI to identify trends and issues that need attention. This means a better experience for consumers, and FIs will be less likely to lose customers or unknowingly engage in practices that harm consumers.

Incident response management. AI can aid in incident response, identifying what laws govern a scenario and providing an appropriate response plan.

Monitoring high-risk customers. FIs bank high-risk customers like money service businesses and cannabis-related businesses. It would be helpful if FIs could use AI to monitor point-of-sale transactions to assess if a transaction looks nefarious. If a high-risk customer were open to sharing internal data with their FIs for AI analysis for fraud and illegal activity, the FI might be able to offer these businesses lower fees or more access to banking.

Compliance. Another risk management application of AI is in analysis and monitoring of regulatory compliance. If AI could parse laws and connect it with and generate relevant KRIs, it could help FIs enhance compliance oversight.

Audit. As audits become more model driven, auditors find themselves looking for the same things over and over. AI could one day automate audits, flagging items that require the attention of human auditors and engaging in hundreds of reviews a week instead of a handful of reviews per month.

Who Controls the Data

An issue that should be top of mind for Congress and the regulatory agencies is how large tech companies that aggregate huge amounts of data use it going forward and the impact it could have on the financial services industry. Just a handful of companies control most of the U.S.'s AI infrastructure. The data troves they hold go far beyond financial information. They know what consumers are searching for, including financial institutions, products, and services.

What happens when one of these companies wants to become a bank or issue credit? They control huge segments of the Internet, the application servers, and the data. What is to stop them from suppressing search results for small FIs or putting their own products and services at the top of search results? They could easily displace the entire banking industry, disadvantaging small FIs in an increasingly digital landscape. What will that do to local communities? What will that do to consumer choice and the competitive landscape?

AI is not going away, and the government needs to be prepared when companies that control the AI infrastructure seek to become a bank. There needs to be strong protections in place further separating banking from commerce (and companies that traffic in data) and the regulatory agencies needs resources to increase their AI sophistication to effectively manage the parties involved in this space.

Ncontracts is grateful for the opportunity to comment on the use of artificial intelligence (AI) tools at financial institutions. If we can provide any further insights or information, please feel free to reach out to us at Stephanie.Lyon@ncontracts.com or Bill.Simpson@ncontracts.com.

Sincerely,

Stephanie Lyon

Vice President, Compliance

Bill Simpson

Chief Product &Technology Officer

Ncontracts
214 Overlook Circle
Brentwood, TN 37027
1-888-370-5552
www.ncontracts.com