Ryan T. Miller
Vice President & Senior Counsel
Innovation Policy
202-663-7675
rmiller@aba.com

August 12, 2024

Mr. Moses Kim
Director, Office of Financial Institutions Policy
Department of the Treasury

Via http://www.regulations.gov/

Re: Docket Number TREAS-DO-2024-0011; Response to Request for Information on Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector

Dear Mr. Kim,

The American Bankers Association (ABA)[1] and 21 state bankers associations representing banks in California, Colorado, Delaware, Hawaii, Iowa, Kentucky, Louisiana, Maryland, Massachusetts, Michigan, Nebraska, Nevada, New York, Ohio, Pennsylvania, South Dakota, Virginia, Washington, West Virginia, Wisconsin, and Wyoming (the Associations) appreciate the opportunity to respond to the request for information (RFI) on the Uses, Opportunities, and Risks of Artificial Intelligence in the Financial Services Sector issued by the Department of the Treasury.[2] Artificial intelligence (AI), including the latest iteration commonly dubbed "generative" AI (GAI), are technological tools that have their place in aiding various banking use cases. However, as with any technology, its deployment should only take place in an environment that carefully considers potential risks with appropriate mechanisms in place to manage those risks. This is especially true for banks, which are integral to the economy, responsible for safeguarding customers' money, and which need to build and maintain customers' trust in order to do so. Moreover, it is becoming increasingly clear that interdicting AI is not a viable option—the technology is too pervasive and promising to block or drive underground.

Banks have employed AI in a responsible manner for decades and are leveraging that mature risk management framework as they begin implementing GAI. Moreover, ABA, the Associations, and our members have had numerous discussions on ways to augment the risk management framework in the wake of GAI, and to that end we present recommendations in the realms of legislation, regulation, and supervisory guidance.

In particular, we wish to highlight two key recommendations at the outset of this comment. First, we urge that any new horizontal federal law pertaining to AI preempt state requirements and clearly exclude banks from any duplicative obligations. As observed in the privacy landscape,

[1] *The American Bankers Association is the voice of the nation's $23.7 trillion banking industry, which is composed of small, regional and large banks that together employ approximately 2.1 million people, safeguard $18.8 trillion in deposits and extend $12.5 trillion in loans.*
[2] https://www.federalregister.gov/documents/2024/06/12/2024-12336/request-for-information-on-uses-opportunities-and-risks-of-artificial-intelligence-in-the-financial#addresses.

the lack of preemption and the confusing applicability to bank data has led to inconsistent levels of consumer protection and significant compliance burden, and policymakers cannot allow this misstep to happen again with AI. Second, we call for updated model risk management guidance from the prudential regulators to clarify expectations in the wake of changes to the ecosystem, but only after an appropriate notice and comment period.

This comment letter is organized in the following manner:

I. Introduction

A. Definitions

For purposes of framing the discussion, ABA members generally agree with the definition used in 15 U.S.C. 9401(3)[3] as well as Treasury's interpretation thereof as noted in the RFI.[4] Some ABA members believe the definition of AI systems[5] by the Organisation for Economic Co-operation and Development (OECD) is superior. These overlapping preferences are reflective of the different internal definitions used by banks in their AI governance programs.

ABA members believe it is necessary to distinguish between traditional AI and GAI. Traditional AI can be thought of as a system designed to respond to a particular set of inputs, such as algorithms and machine learning. It "learns" from the data and makes decisions or predictions based upon the data, but does not create anything new.

For GAI, ABA members are supportive of the definition used in President Biden's Executive Order 14110: "the class of AI models that emulate the structure and characteristics of input data

[3] A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human--based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

[4] Describing a wide range of models and tools that utilize data, patterns, and other informational inputs to generate outputs—including statistical relationships, forecasts, content, and recommendations—for a given set of objectives. For the purposes of this RFI, Treasury is seeking comment on the latest developments in AI technologies and applications, including but not limited to advancements in existing AI (*e.g.*, machine learning models that learn from data and automatically adapt and improve with minimal human interference, rather than relying on explicit programming) and emerging AI technologies including deep learning neutral network such as generative AI and large language models.

[5] An AI system is a machine-based system 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. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

in order to generate derived synthetic content...[t]his can include images, videos, audio, text, and other digital content."[6]

Definitions of AI, AI systems, and GAI will continue to evolve and consensus will settle in as more entities build and mature their internal AI governance programs, but it is important to bear in mind the conceptual differences between different iterations of the technology.

However, it is one thing to agree on a definition to form the basis of a conversation, quite another to use in the context of laws and regulations. As noted above, AI and GAI are tools that are utilized in the context of specific activities. Therefore, any requirements should be rooted in those activities and the risks presented, not in the technology that helps deliver them. AI and GAI fall into a spectrum and it is not always a simple matter to categorize given software as AI or another type of computer program.

ABA and the Associations strongly believe that any new laws and regulations applying to AI activity must be industry-focused, risk-based, and tied to use case. Because variations in definitions across agencies can cause confusion within individual banks and in the ecosystem at large, we suggest the National Institute of Standards and Technology (NIST) work with banking regulators to craft a workable interagency definition that can be leveraged to guide specific policies, particularly in the realm of supervisory activity.

B. Current State

It is imperative to understand the existing strong culture of compliance in which banks operate: three lines of internal risk-management defenses (see pages 11-12 for more information on this concept); application of technology-neutral laws, regulations, and guidance; and validation of the effectiveness of the framework through regular examinations by the bank regulatory agencies.

Banks have long used AI, although traditional AI is far more prevalent than emerging GAI today. This deployment has been subject to oversight by prudential regulators, the Consumer Financial Protection Bureau (CFPB), and other bodies. For example, the Federal Reserve, Office of the Comptroller of the Currency (OCC), and the Federal Deposit Insurance Corporation (FDIC) have all adopted supervisory guidance on model risk management.[7] The guidance applies to models, which are defined as:

> a quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. A *model* consists of three components: an information input component, which delivers assumptions and data to the model; a processing component, which transforms inputs into estimates; and a reporting component, which translates the estimates into

[6] https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

[7] SR 11-7, OCC Bulletin 2011-12, and FIL-22-2017, respectively. The OCC also released a booklet for its examiners to use as an aid when supervising banks' model risk management programs; *see* https://www.occ.treas.gov/publications-and-resources/publications/comptrollers-handbook/files/model-risk-management/index-model-risk-management.html.

> useful business information…The definition of *model* also covers quantitative approaches whose inputs are partially or wholly qualitative or based on expert judgment, provided that the output is quantitative in nature."[8]

Further, even "more qualitative approaches used by banking organizations— i.e., those not defined as models according to this guidance—should also be subject to a rigorous control process."[9] Thus, treatment should be scaled based on the use case and risk appetite. In 2021, the Federal Reserve, the OCC, and the FDIC issued interagency guidance addressing model risk management to support Bank Secrecy Act/Anti-Money Laundering and Office of Foreign Assets Control Compliance (BSA/AML and OFAC).[10]

Interagency guidance issued by the Federal Reserve, the OCC, and the FDIC on third-party risk management is crucial to navigating the AI ecosystem.[11] The document is principles-based and technology-neutral, which is appropriate given the wide range of third-party relationships in which banks engage and the diversity in bank size and complexity in the financial services ecosystem. This optimizes the ability of banks to identify concerns germane to their business model and the purposes for which the third party's technology will be used as they conduct due diligence and develop appropriate risk mitigants for third-party relationships.

Several agencies have clarified that consumer protection and anti-discrimination laws continue to apply whether or not AI or GAI is utilized.[12] The CFPB released guidance pointing out risks of AI present in chatbots[13] as well as the importance of explainability in complying with Regulation B.[14] Banks are well-positioned to mitigate the potential for discrimination in AI through their robust risk management and compliance management systems. However, because the use of opaque models can complicate compliance with requirements to provide reasons for adverse action to credit applicants, the industry needs clear and consistent guidance. To that end, the CFPB should reconcile recently issued guidance on adverse action and models with its existing official interpretation of the Equal Credit Opportunity Act (ECOA).[15]

---

[8] SR 11-7, page 3.
[9] Id.
[10] SR 21-8, OCC Bulletin 2021-19, and FIL-27-2021, respectively.
[11] Guidance on Third-Party Relationships: Risk Management, https://www.federalregister.gov/documents/2023/06/09/2023-12340/interagency-guidance-on-third-party-relationships-risk-management.
[12] Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems, https://files.consumerfinance.gov/f/documents/cfpb_joint-statement-enforcement-against-discrimination-bias-automated-systems_2023-04.pdf.
[13] CFPB, Advisory on chatbots in consumer finance, https://www.consumerfinance.gov/data-research/research-reports/chatbots-in-consumer-finance/chatbots-in-consumer-finance/.
[14] CFPB, Circular 2023-03, "Adverse action notification requirements and the proper use of the CFPB's sample forms provided in regulation B," https://files.consumerfinance.gov/f/documents/cfpb_adverse_action_notice_circular_2023-09.pdf.
[15] Id.

Any clarifying guidance on AI usage issued by agencies must be published in advance for stakeholder feedback.[16] Further, as agencies develop this guidance they should be mindful of the challenges faced by small banks looking to AI for competitive reasons but lacking deep benches of internal experts.

Treasury and the Department of Homeland Security have issued reports that may help banks and other financial institutions identify, assess, and mitigate certain forms of risk presented by AI-enabled use cases.[17] The Treasury report included a paper produced by the Financial Services Sector Coordinating Council's R&D Committee (FSSCC R&D), which ABA co-chairs.[18]

FS-ISAC, a global non-profit that advances cybersecurity and resilience in the global financial system, has issued six papers that aim to help the financial services sector capitalize on AI's opportunities while mitigating its risks.[19] Specifically, the papers entitled "Adversarial AI Frameworks: Taxonomy, Threat Landscape and Control Frameworks"[20] and "Responsible AI Principles"[21] define several key terms and lay out principles that may be of assistance to financial institutions building up or maturing their AI governance.

Finally, NIST has developed a voluntary, industry-agnostic, and customizable AI Risk Management Framework (AI RMF) that banks may use to develop and assess their AI governance programs.[22] NIST recently unveiled its initial public draft of a GAI supplement to the framework.[23] In a comment letter, ABA expressed support for the initial public draft and made several recommendations on how it could be improved.[24]

Together, these documents aid the creation and maturation of banks' policies, procedures, and governance regarding new technologies (including AI and GAI). Banks continuously monitor and update these programs as risks evolve.

As a result of this careful approach, use of GAI by banks is still in its nascent stages. Non-banks offering financial services, by contrast, are not subject to many of these obligations and thus may be less mindful of the risks before offering GAI. We caution policymakers to ensure that any

[16] ABA also issued a whitepaper, Effective Agency Guidance, recommending steps agencies could take to issue guidance that complies with agencies' legal requirements while providing useful advice and information to regulated entities. *See* https://www.aba.com/advocacy/policy-analysis/wp-effective-agency-guidance.
[17] *See* Treasury Report on Managing Artificial Intelligence-Specific Cybersecurity Risks in the Financial Services Sector, https://home.treasury.gov/news/press-releases/jy2212; *see also* DHS Report on Mitigating AI Risk: Safety and Security Guidelines for Critical Infrastructure Owners and Operators, https://www.dhs.gov/publication/safety-and-security-guidelines-critical-infrastructure-owners-and-operators.
[18] https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf.
[19] https://www.fsisac.com/knowledge/ai-risk.
[20] https://www.fsisac.com/hubfs/Knowledge/AI/FSISAC_Adversarial-AI-Framework-TaxonomyThreatLandscapeAndControlFrameworks.pdf.
[21] https://www.fsisac.com/hubfs/Knowledge/AI/FSISAC_ResponsibleAI-Principles.pdf.
[22] NIST AI RFM, https://airc.nist.gov/AI_RMF_Knowledge_Base/AI_RMF.
[23] NIST AI 600-1, https://airc.nist.gov/docs/NIST.AI.600-1.GenAI-Profile.ipd.pdf.
[24] *See* https://www.regulations.gov/comment/NIST-2024-0001-0154.

new laws, regulations, or guidance are consistent with the above framework already followed by banks to such good effect.

II. General Uses of AI

A. Forms of AI

Banks have a long history of using traditional AI within the risk management framework outlined above. Traditional AI is generally predictable and has been vetted through years of usage and supervisory feedback. Typical traditional AI use cases include fraud detection and prevention, marketing, cybersecurity, AML activity, credit underwriting, and customer service.

As noted above, GAI applications are still in the nascent phase; these are still early days when it comes to developing trust in the technology and understanding regulators' expectations. Banks are aware of the need to identify the right stakeholders to build proper guardrails and are proceeding cautiously. If GAI applications become more common at banks, regulators can be confident that those banks have carefully considered and managed the associated risks because of banks' requirement to apply comprehensive risk management processes to GAI or any other new technology. By comparison, non-banks offering similar financial services are frequently not required to, and often do not, apply the same risk management controls.

Banks are focused on effective risk management, governance, and controls when engaging GAI. This includes pursuing approaches to further enhance data protections, improve accuracy, and minimize hallucination. ABA members continue to investigate the appropriate use for large language models (LLMs) trained on vast quantities of data as well as small language models that are trained on first party data. In addition, banks are investigating purpose-built models for narrow applications in addition to models designed for general use.

B. Reliance on Third-Party Providers

Banks of all sizes use third-party sources to develop AI models, and the challenges associated with managing the risk posed by these models are universal across the banking industry. Moreover, there is not a level playing field between highly regulated banks and comparatively unregulated Big Tech developers. This disparity illustrates the need for a standard-setting organization to help address transparency and model validation challenges, which we discuss further in Sections IV(D) and V of this letter.

Many community banks purchase "off the shelf" products because they do not have the technical staff or other resources to build and maintain their own models, or require greater stores of training data than they would otherwise have access to. However, some community banks demonstrate more innovative behavior due to especially supportive leadership or forward-leaning personnel.

Regional and large banks also utilize third-party AI, although it is common for these institutions to fine-tune the models themselves or request that the third party customize them. However, the number of third-party developers of GAI foundational models is small, which markedly limits

the ability of even the largest banks to obtain bespoke products. Large banks also develop some models in house. In short, large banks have options in AI development and deployment that most community banks do not.

Banks of all sizes encounter due diligence and oversight challenges associated with third-party models. Vendors rarely assist with bank efforts to validate the third-party product and may share necessary details only begrudgingly notwithstanding the increasing complexity of models. Banks attempt to include contractual language requiring third parties to be transparent, but vendors often decline to do so in an effort to protect their proprietary information.

Further, community banks often lack the market power to negotiate these types of provisions, and even large banks experience challenges bargaining with the Big Tech developers. Banks of all sizes also encounter regulatory headwinds in their development and deployment of AI as examiners have created confusion regarding the expected level of AI documentation. A specific example encountered by community bankers was examiners expecting personnel to explain the software code that powered AI models, rather than focusing on truly important matters such as the model inputs, outputs, and outcomes. Moreover, the complexity of LLMs is such that even global systemically important banks (GSIBs) may not be able to address things on a granular level. This emphasizes the need to focus on the levers that can actually have an impact.[25]

C. Use Cases

The below use cases predominantly utilize traditional forms of AI unless otherwise noted. As previously stressed, GAI is not deployed at scale by banks at the time of this writing.

Cybersecurity

AI is used to detect and respond to potential cyberattacks more quickly and efficiently than human intelligence could accomplish alone. AI-based network vulnerability security software and services can detect and monitor incoming and outgoing network traffic to identify suspicious patterns to aid security analysts in their initial detection and classifications. While AI is designed to scale up the analysts' work by reducing the time spent on false positives, the role of human verification is essential at this time. By involving human oversight (often referred to as "human in the loop"), potential errors are avoided, thereby striking the proper balance between automated efficiency and operational safety. As AI technology advances, evolving cybersecurity risk

[25] For more information on the concentration risk theme, *see* the **Financial Sector Cloud Outsourcing Issues and Considerations** document which seeks to address challenges raised in the Treasury Cloud Report related to transparency, resource gaps, exposure to operational incidents originating at cloud service providers (CSPs) and contract negotiation dynamics. The document, authored collectively by the FSSCC Cloud Outsourcing Issues and Considerations Workstream and the American Bankers Association (ABA) with support from the Securities Industry and Financial Markets Association (SIFMA), identifies a non-exhaustive list of key considerations for developing contractual provisions between financial institutions and CSPs to address risks, regulatory and supervisory compliance expectations when using cloud services. These key considerations should be used as a voluntary reference tool by financial institutions during the contract negotiation phase of onboarding a CSP to appropriately address cybersecurity, resilience, and third party-due diligence expectations, and to enable compliance with growing financial services regulatory requirements and supervisory expectations. The document is available at: https://www.aba.com/news-research/analysis-guides/fsscc-cloud-outsourcing-issues-and-considerations-july-2024.

management practices to effectively detect and mitigate these complex and emerging risks is crucial. In addition, banks may also utilize GAI to pinpoint malicious code as well as aide internal developers in identifying vulnerabilities in their own code.

Fraud Detection & Prevention

AI models using predictive analytics help banks proactively find anomalies in transactions and identify outliers that do not conform with customers' past patterns or payment activity. Financial institutions work with payments companies to leverage AI to help mitigate fraud and promote smarter authorization, clearing, and settlement for card-based and non-card-based payments. These partnerships are especially helpful in the context of digital payments, and use of traditional AI in this space goes back multiple decades. AI thus supports fraud detection, which can minimize losses and preclude the need for remedial action. In other words, AI models not only improve the performance of fraud detection capabilities, but also help catch fraudulent activity before it impacts customers.

Lending

Banks use AI across lending processes to help identify accounts that can be approved for credit, as well as loan amounts and pricing. AI thereby assists banks with evaluation of creditworthiness and improves efficiency in decisioning. Another lending use case is providing metrics around key life indicators such as attrition rates for mortgages. We emphasize that banks are taking a very cautious approach with the integration of GAI into lending and credit underwriting decisions.

Customer Service

AI assists banks with learning how customers are interacting with products and services. In addition, AI can perform sentiment analysis to gain insight into satisfaction. This data can be fed into a platform to better understand customer interactions and how to improve them. It is anticipated that GAI will streamline this use case given its ability to summarize conversations with contact center personnel and may also be used to help representatives grasp the backstory when customers have to re-engage on the same issue with a different individual. Customer service is an early area for GAI focus as it can increase the quality of support and enables agents to be more efficient, while the agent serves as a "human in the loop" control function.

Chatbots

Chatbots are a form of customer service but warrant their own heading. Chatbots powered by traditional AI are commonly used and can respond with static responses to certain keywords. Customers often gravitate towards these channels due to ease of use and preference for self-service. The CFPB has expressed considerable reservations around the use of chatbots and issued an advisory to this effect.[26] Of course, chatbots that provide incorrect information can cause

[26] Supra, note 13.

harm to consumers (see, for example, the case of the Air Canada chatbot).[27] However, imperfect information is not an issue unique to AI, and as is the case with human employees can be remediated with necessary training. Moreover, records from deployed chatbot conversations can be used to improve future performance. Banks run the risk of reputational harm if chatbots exhibit faulty functionality, and for that reason conduct extensive vetting of "tuning" so as not to enter risky areas such as negotiation of rates or giving financial advice.

And yet, banks are firm in the potential virtues of chatbots powered by AI and GAI. Not all will choose to deploy a GAI chatbot, but it is essential that their ability to do so in a safe and responsible manner be preserved. Routine questions could be handled by the chatbot, with referral to human beings for more complex issues. Functionality will continue to improve as time goes on, allowing the chatbot to handle more situations. With the proper controls (vetting and validation of the knowledge sources, guidance, guardrails and continuous feedback loops), chatbots can be a key component of banks' customer service and employee training.

Marketing

GAI can be used to automate the creation of various types of marketing content, including ad copy, social media posts, image/video generation, and product descriptions. AI can personalize content to optimize its relevance. Personalization must be balanced with privacy principles that protect personal information and respect consent. Accuracy, purpose limitations, and data/storage minimization are key privacy pillars along with transparent communications with the data subject. Any content creation assisted by AI should follow the same review and approval path as human-created materials.

Back-Office

There are several back-office use cases that banks are exploring:

Coding- Banks of all sizes have internal development teams experimenting with the use of GAI to generate code in a variety of use cases. GAI's ability to assist with routine coding is an area that shows promise when complemented by a human coder for oversight. GAI can increase the productivity of the human coder. Like other technology areas, free and open-source tools may be used in AI/GAI development to improve productivity in creating GAI prototypes and solutions. However, the process must follow appropriate software development procedures (including testing) to protect against attempted malware injection.

Regulatory Reporting- GAI can assist with compiling the necessary data points for completing regulatory reporting. Of course, this would have to be validated prior to submission to confirm accuracy.

Internal Document/Knowledge Management- An intriguing use case is connecting bank resources into an internal document management system. There are tools to gather necessary

[27] *See* https://www.bbc.com/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know.

information across disparate sources and compile that information for employees (for example, contact center representatives). The nature of GAI makes it easier to ingest policies and procedures, which could allow for harmonization and identification of gaps in historically siloed departments.[28] GAI may bring significant improvements in the use of content management systems due to its ability to query and find related information to answer a question that crosses over multiple documents and multiple document types.

Second & Third Line- Banks can potentially use GAI to generate a first draft of policies and procedures, or to summarize a description of applicable controls. It can also aid in creating uniform formats that can make it easier to comprehend information submitted from various sources (for example, vendor information for due diligence purposes or the creation of SAR narratives). Once these improved policies and procedures are created, they can be ingested into the GAI RAG stores[29] to provide operational efficiencies in locating a desired policy or procedure.

III. Actual and Potential Opportunities

The integration of AI and GAI offers the financial services sector increased efficiency, precision, and adaptability, as well as the potential to bolster the resiliency of institutions' systems, data, and services. This opportunity is set against the reality of a complex and persistent threat landscape that is also using AI for malicious purposes.

Increased efficiency seems to be the most commonly referenced opportunity associated with AI and GAI. In particular, BSA/AML and fraud applications were mentioned as areas where efficiency most frequently manifests due to superior pattern recognition (which will only improve with time). However, efficiency and productivity are not simply banking issues but can help employees in many sectors automate manual tasks and free up bandwidth. Banks should be able to avail themselves of this opportunity in a responsible way, just as other industries can. In addition, AI can help improve operational resiliency in areas such as transaction authorizations by having backup systems come online in the event of outages caused by natural disasters or other occurrences.

Another potential benefit of GAI is improved credit decisioning, which would allow lending to consumers who previously might have been denied. This might be due to ingestion of other, alternative forms of data not previously usable. However, this is still theoretical and has not been borne out due to the cautious deployment for lending applications.

It is also possible that increased use of AI and GAI will improve customer service; for example, through the use of chatbots and virtual assistants, which can support the effort of live agents and provide training data to increase performance. Compliance validation at banks is sufficient to

[28] However, at present GAI will not solve the problem of incorrect or inconsistent source data. Before connection, source data must be extracted and ingested into Retrieval Augmented Generation (RAG) data stores that structures prompts with the most relevant bank knowledge sources. The perfection of these RAG processes will highlight and identify where the source data needs to be corrected or contextualized to yield optimal responses. RAG processes can be developed so that the relevant data can be conversationally queried and answered in an iterative fashion.
[29] *See* supra, note 28.

ensure deployed programs comply with consumer protection laws, are accurate, and are capable of capturing complaints for further action.

IV. Managing Actual and Potential Risks

In the financial sector, risk management methodologies are pivotal in maintaining integrity and stability, and banks (though not necessarily non-banks offering similar services) are required to implement appropriate risk management frameworks. The three lines of defense system serves as a foundation for banks, promoting rigorous oversight and clear delineation of responsibilities among operational management, risk/legal/compliance, and internal audit functions. These processes are created using supervisory guidance with feedback from regulators. They are technology-neutral and risk-based, and are constantly updated to reflect emerging risks and conform to prevailing best practices. This philosophy allows banks to achieve practical results without getting bogged down in definitions; for example, in preparing model inventories. Banks are in the best position to apply the spirit of the requirements and react to inputs from their regulators—a symbiotic process that has led to positive results time and again.

There should not be a separate workflow for evaluating AI applications as such; instead, the existing risk management framework should become "AI aware" to flag the inherent risk associated with AI-enabled use cases. A material feature of GAI to bear in mind is its prompt-based nature, which lowers the barrier to access and expands the pool of potential users. Banks have been breaking down silos in favor of enterprise-wide functions, and this movement will continue. As part of this transition, banks have been building interdisciplinary teams to identify, assess, and mitigate risks stemming from particular use cases. These cross-functional groups allow for superior risk awareness and mitigation versus compartmentalized approaches.

Many ABA members have been using AI technologies for years, while being subject to supervision and regulation. Their existing and mature governance frameworks are already mitigating AI-specific risks. Further, ABA members are confident that banks will be able to adapt their existing risk management and governance frameworks to mitigate risks associated with increased use of AI and with emerging GAI technologies. This is testament to the three lines of defense structure used by banks coupled with agency supervision. The gap is how entities without such supervision, such as fintechs,[30] will fare with respect to safeguarding consumers and preserving financial stability.

The "three lines of defense" refers to the division of roles and responsibilities within a bank in order to identify, assess, and mitigate risks. Among other sources, this expectation is provided in the model risk management expectations guidance, SR 11-7:[31]

[30] Treasury sought comment from a broad array of stakeholders. The RFI defined "financial institutions" as "any company that facilitates or provides financial products or services…includ[ing] banks, credit unions, insurance companies, non-bank financial companies, financial technology companies (also known as fintech companies), asset managers, broker-dealers, investment advisors, other securities and derivatives markets participants or intermediaries, money transmitters, and any other company that facilitates or provides financial products or services under the regulatory authority of the federal financial regulators and state financial or securities regulators." *See* supra, note 2.

[31] SR 11-7, see pp. 18-19.

- The first line are business units, which are generally responsible for the risk associated with their business strategies. They are ultimately accountable for the risk and performance within the framework set by bank policies and procedures, and are responsible for ensuring processes are properly developed, used, and evaluated.

- The second line is the control function. The responsibilities include risk measurement, limits, and monitoring. Other responsibilities include managing the independent validation and review process to ensure that effective challenge takes place. Control staff should have the authority to restrict business operations and order corrective action. Control work can be done in a way that prioritizes the greatest risk.

- The third line is the bank's internal audit function. The third line's role is not to duplicate risk management activities but to evaluate whether risk management is comprehensive, rigorous, and effective. They should be independent and document findings. The third line should possess expertise but should not be involved in the first or second line of work. The third line should also verify that acceptable policies are in place, owners and control groups comply with those policies, validation work is conducted properly, and appropriate degrees of effective challenge is being carried out.

We now turn to a discussion of particular risks presented by AI and GAI applications that banks mitigate through their governance structure.

A. Cybersecurity & Fraud

One of the sector's foremost concerns is the acceleration of threat actors' capabilities due to AI, particularly GAI. As financial services becomes increasingly digital, they become ever more vulnerable to cyber attacks. While all AI technologies enable threat actors to deploy advanced attack tactics, the challenge with GAI lies in the reduced cycle time for these actors. Skilled adversaries aided by GAI can swiftly identify and implement novel breach techniques, potentially outpacing traditional detection strategies. This necessitates a continuous and rapid update in detection methodologies to address cyber threats that financial institutions might face. While developers are implementing guardrails to deter this activity, security researchers have demonstrated that it is not difficult to circumvent the guardrails, necessitating robust, layered defenses. Current capabilities may not be fully equipped to address these novel threats, necessitating enhancements in both technical capabilities and control processes.

As GAI tools become more accessible and sophisticated, they are increasingly exploited by bad actors and mingled with social engineering to perfect phishing messages, clone voices, and simulate video conferences. GAI can also be used to create advanced malware. These attacks are proceeding at an alarming pace and have the potential to bypass legacy detection methods.

As a result, banks need to play both defense and offense in order to combat the bad actors. This requires education and familiarity with their systems, as well as technical tools. Banks currently have processes to rapidly respond to and patch known vulnerabilities. The focus is on earlier

identification, where AI has had and will continue to play a major role. In addition, processes such as red teaming are being expanded to address earlier identification of GAI threats.

In discussing the adversarial use of AI in cybercrime, it is important to differentiate between key concepts: adversarial machine learning,[32] adversarial training,[33] and adversarial attack learning.[34] These concepts are distinct from the use of GAI for generating credible-looking text, images, code, and malware, which often dominate headlines but are not the only threat.

The evolution of GAI content and deepfake creation services are of concern to the financial sector, especially smaller and less well-resourced institutions. These technologies not only lower the barrier to entry but also complicate authenticity verification measures. Bad actors can leverage tools to fool employees and customers. In addition, GAI can produce authentic-looking documents that can be further used to establish synthetic/fake identities, impersonate true persons, engage in damaging brand and reputational conduct, or gain access to systems.

ABA and the Associations appreciate Treasury's efforts to collaborate with the financial services sector via the public/private collaboration under FSSCC and the Federal Banking Information Infrastructure Committee (FBIIC). One of the projects under development is the formation of best practices for financial institutions and technology companies to mitigate identity-related risks tied to AI-generated impersonations (i.e., deepfakes). As part of the project, ABA is working to identify steps government can take to close the gap between physical and digital government credentials – as well as enabling identity information to be validated against government "reservoirs of truth" – to enhance the digital identity ecosystem that financial institutions rely on for safe, secure, privacy-preserving, and reliable transactions. ABA and the Associations anticipate this initiative to improve the ability of financial institutions to trust and to digitally validate government ID documents and other digital attributes during enrollment and authentication processes.

B. Privacy & Data Governance

We have increasingly seen policymakers around the world recognize the important dovetail between data privacy, AI policy, and risk management issues. The training phase of GAI presents privacy risk around the scope of the data collected, the duration of data retention, whether it is clear the data is being used to train AI models, and whether the data subject

[32] Adversarial Machine Learning is a field of study that focuses on the security of machine learning systems. It involves understanding how attackers can manipulate or exploit machine learning models and developing techniques to defend against such attacks. This is a technique in machine learning where the model is trained to make decisions, and an adversary intentionally inputs deceptive data to cause the model to make a wrong decision. Banks must therefore test the robustness of the model against intentional manipulation.

[33] Adversarial Training is a defense mechanism against adversarial attacks. It involves training the machine learning model on a mixture of clean and adversarial examples, which are designed to be difficult for the model to classify. The goal is to make the model more robust against manipulation by exposing it to a variety of attack strategies during training.

[34] Adversarial Attack Learning refers to the process of learning how to perform adversarial attacks on machine learning models. It involves understanding the model's weaknesses and developing methods to exploit them, often by creating input data that the model will misclassify. Concerns also extend to prompt injections into various forms of LLMs, with the speed of patching varying depending on deployment methods.

consents to that usage. Banks are working to implement safeguards to ensure that the GAI only uses personal information consistent with disclosures and expectations.

Perhaps even more critical is the risk to banks' sensitive data (including customer personal information) to prevent data leakage. RAG techniques[35] with large foundational models allow banks to protect their proprietary input from being exposed to any LLM learning or training. This has the added benefit of confirming the data's accuracy and that it yields the desired model output.

Data provenance, quality, and permissibility are among the biggest risks encountered in AI and GAI. This is true both for traditional data historically used and gathered by banks for modeling purposes, as well as non-traditional data such as that originally collected for other purposes or obtained from reputable third-party sources. ABA members anticipate state legislative activity in these areas could further complicate the situation. A potential mitigating factor is the use of synthetic data—which can be derived from real data and conceal actual personal information. This might then in turn be used to train other models and may assuage concerns over handling sensitive data within AI applications.

ABA members highlighted the criticality of data governance programs to ensure the right permissions are associated with the training data. Effective governance can only occur when the proper stakeholders are included in the associated committees (namely, business units, IT staff, legal, compliance, information security, risk officers, among others).

Enterprise privacy programs already exist at many banks and can serve as a foundation for the broader concept of AI governance. Moreover, banks pursue specific use cases and manage risk as they always have, whether AI is involved or not. There is not a special path; rather, concepts around managing risk of AI are integrated into the existing framework.

There are also open questions about forthcoming requirements such as Section 1033/personal financial data rights[36] and how data shared pursuant to customer consent can be used for training AI models, both internally and at third parties (including service providers). Excessive controls placed around data used to train AI can fritter away potential benefits. Nonetheless, adequate mitigants are necessary because the stakes associated with data loss/leakage are higher in a world where GAI can produce digital clones and construct synthetic identities.

A known issue with GAI tools is that they can "hallucinate" or create incorrect information that appears legitimate. This phenomenon underscores the need to have mechanisms in place to instruct the model not to respond if it lacks sufficient data.

[35] *See* supra, note 28.
[36] https://www.federalregister.gov/documents/2023/10/31/2023-23576/required-rulemaking-on-personal-financial-data-rights.

C. Bias & Fair Lending

AI holds promise in making financial services and products more broadly available. For example, models' use of rental payment history can expand access to credit for consumers who lack traditional credit scores. Banks are excited about the possibility of serving credit needs of these and other underserved populations. At the same time, banks understand that they must consider the fair lending implications of AI use, and that there is no "AI exemption" from consumer protection and fair lending laws. As addressed above, regulators have made it clear that consumer protection and anti-discrimination laws apply to AI-powered use cases.[37]

Fair lending risks take the form of disparate treatment, which could result from a model's inclusion of attributes (or their proxies) prohibited by ECOA and the Fair Housing Act (FHA). This includes disparate impact, which results from the unjustified use of neutral factors that disproportionately affect applicants on prohibited bases. Fair lending also requires that model outputs be explainable to applicants, as ECOA requires a creditor to provide the principal reason(s) for denying an application for credit.

Banks report that their existing risk governance regimes, including model review, have generally helped them identify and mitigate the fair lending risks in using AI. However, it can be challenging to identify variables and their interaction in complex models, and to test models for disparate impact. Collaboration between the bank's lines of business, modelers, analysts, and fair lending officers is critical to these efforts.

A flexible approach to regulation in this space is important to allow for the evolution of AI and to recognize that banks have different levels of complexity and risk. We note that a flexible approach was taken in a recent six-agency rule on quality control standards for automated valuation models used in mortgage originations and securitization determinations.[38] It is useful to consider that while the agencies included compliance with fair lending laws as a component of quality control in AVM use, they decided not to prescribe a specific means to achieve quality control. They noted that their "flexible approach to implementing the quality control standards provided by the final rule will allow the implementation of the standards to evolve along with changes in AVM technology and minimize compliance costs." They further stated that "[i]nstitutions will have flexibility to adopt approaches to implement this quality control factor in ways that reflect the risks and complexities of their individual business models."[39]

[37] *See* supra, notes 12-14.

[38] The final rule implements section 1125 of the Financial Institutions Reform, Recovery, and Enforcement Act of 1989 (FIRREA), as amended by the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank Act), which directs the agencies to require quality control standards for AVMs used in connection with making credit decisions or in covered securitization determinations for a mortgage/mortgage-backed security. The final rule requires that AVMs have sufficient quality control standards to: (1) ensure a high level of confidence in the estimates produced; (2) protect against manipulation of the data; (3) seek to avoid conflicts of interest; (4) require random sample testing and reviews; and (5) comply with nondiscrimination laws. In adding the fifth requirement, the agencies cited the need to "raise awareness of" applicable fair lending laws in AVMs. *See* https://www.federalregister.gov/documents/2024/08/07/2024-16197/quality-control-standards-for-automated-valuation-models.

[39] Id.

Explainability is important in financial services applications of AI as it can help ensure that model outputs or decisions are not biased, assist developers in improving models, and increase user confidence in the outputs of the model. But AI may generate outputs where the basis is difficult or impossible to determine. Practices around data input, decision-making criteria and weighting of those criteria, assurance review and others are being developed to ensure that validation processes keep pace with technology. In addition, banks are tracing how AI models process inputs into outputs to better understand the states of the models before and after processing. Financial institutions already build AI functionality with explainability in mind. Risk management practices in the financial sector are mature and include incorporating relevant elements from the NIST AI RMF.

While banks have sound processes in place to address fair lending risks and explainability, and ECOA and FHA provide a robust framework to address bias and discrimination in AI, additional guidance in a few areas could help banks meet their obligations to serve all customers fairly.

First, smaller banks may need to leverage AI for efficiency but many lack the resources to validate and test AI. Assessing and addressing disparate impact risk stemming from AI can be a complicated, lengthy, and expensive process, given the complexity of new models and the sheer amount of data that can be manipulated. Data attributes may be bundled and cannot be readily separated, vendors may refuse to share attributes, or they may refuse to validate predictability if the bank wants certain attributes removed. Any of these eventualities may force the bank to cease doing business with the vendor. As discussed below, the agencies should clarify expectations around third-party models, including information that they must provide to banks.

Second, for explainability and adverse action notices, regulators should provide clear and consistent guidance. The CFPB recently published a circular on adverse action notices and AI that may be inconsistent with the Bureau's official interpretation of Regulation B.[40] The CFPB's recent circular declares that a creditor must not only disclose the factor that resulted in adverse action, but must also provide more specificity about the factor. For example, according to the circular, it would be insufficient for the creditor to state "purchasing history" or "disfavored business patronage" as the principal reason for adverse action, without more detail, such as the business location, the type of goods purchased, and so on.[41] In contrast, the official commentary states that a creditor may provide a reason such as "age of collateral" even if such a factor's relationship to credit worthiness may not be clear to the applicant.[42] The CFPB's varying statements on a creditor's responsibility have resulted in confusion and may discourage creditors

[40] *See* supra, note 14.
[41] Id.
[42] 12 CFR part 1002, Supp. I, Comments 9(b)(2)-3 and 9(b)(2)-4. In addition, in 2020, the CFPB reaffirmed the official commentary to Regulation B, noting that it could be helpful to creditors using AI. However, the CFPB has labeled this statement "incomplete." The 2020 statement says: "This flexibility [in the commentary] may be useful to creditors when issuing adverse action notices based on AI models where the variables and key reasons are known, but which may rely upon non-intuitive relationships." Innovation spotlight: Providing adverse action notices when using AI/ML models, July 7, 2020, available at https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/.

from beneficial use of AI.[43] The CFPB should clarify its expectations and should coordinate with other regulators so that there is a level playing field for all creditors.

Finally, the proliferation of biometric privacy by laws by the states could lead to difficulties in offering insurance coverage due to poor data quality. This could impact banks due to close ties to the insurance industry in the form of affiliates and joint marketing partnerships.

D. Third-Party Risk & Solutions

Banks routinely rely on third parties for a range of products, services, and other activities. Today, third parties utilize traditional forms of AI to enhance their products and services and improve the efficiency of their internal processes. Banks have extensive policies, procedures, and controls in place to manage the risks associated with these types of AI deployed/provided by third parties. By contrast, third party use of GAI is in its infancy, but is evolving rapidly and will likely become more widespread. This poses both opportunities and challenges for banks, who must manage the operational, compliance, and strategic risks associated with third parties that use AI.

AI holds both promises—including increased efficiency, expanded access to human capital, new delivery channels—and pitfalls, such as heightened risk for privacy violations and fair lending issues. As such, it is important that banks deploy AI in an intentional and responsible manner. The need for a methodical approach to AI also extends to a bank's third-party relationships, which could expose a bank to unknowingly utilizing GAI from a third party and taking on the associated risks.

As we discuss more fully below, some banks are taking steps to ensure that third parties are not utilizing AI without the bank's knowledge and consent. These institutions are also reviewing and revising their third-party risk management programs to take GAI into account. However, many community banks have not yet evaluated whether and to what extent their third parties are using AI. These limitations are often due to a dearth of model and technological expertise, staffing constraints, the increasing complexity of third-party models, and the need to focus on a wide range of new regulatory demands. Large banks also face significant difficulty in this area due to the lack of a level playing field among banks and non-banks.

Examiners are also applying varying levels of scrutiny to AI use by third parties. Some examiners expect banks to provide a detailed explanation as to how their third parties are using AI, while other examiners simply inquire how banks have been approaching it.

Below we explain the key challenges stemming from third-party development and deployment of AI. We also provide recommendations for addressing AI-related challenges and discuss how banks' third-party risk management programs are evolving to account for the risks associated with third parties that use AI.

[43] In response, ABA called for the CFPB to rescind the circular and clarify the requirements for adverse action. https://www.aba.com/advocacy/policy-analysis/letter-to-cfpb-on-ecoa-circular.

Banks Must Have Visibility as to Whether Third Parties Utilize AI for Prioritized Risk Areas

Third party risk management involves the following processes: planning and scoping, due diligence and third-party selection, contract negotiation, ongoing monitoring, and termination. For banks to manage AI risk at each of these stages of the third-party lifecycle, it is incumbent on third parties to disclose whether and how they use AI and to provide visibility into it. Third parties should also be required to inform banks if they begin to use AI during the course of an existing relationship. This information is fundamental to a bank's ability to evaluate a third-party's products and services, determine whether it is willing to accept the inherent risk of that relationship, and ascertain whether it has the capacity and expertise manage the AI risk accordingly. This is especially important with GAI applications.

Even when a third party discloses that it uses AI, banks may not be able to opt out of these technological features. Some large banks have been successful in requesting customization of third-party products, but community banks often do not have the market power to negotiate these terms. And, as a practical matter, some third-party products and services may not function as advertised if the AI features are removed. As AI becomes more widely adopted and is embedded into more vendor processes and products (whether through the bank's third party or elsewhere in the supply chain), we anticipate that banks of all sizes will find it ever more difficult to negotiate or control exposure to third-party AI risk.

A Standard Setting Organization May Help to Address Visibility Issues

Over the years, many fintechs and other technology firms have been reluctant or unwilling to provide visibility into their proprietary models or assist with model validation. Banks report significant variability in the depth of information that third parties provide regarding their model testing procedures and in how they address potential biases in AI models and associated products. In part, this lack of cooperation has been driven by the desire to protect the firm's intellectual property and by an insufficient understanding of the breadth and depth of laws and regulations to which banks are subject. The potential economic reward associated with AI exacerbates these challenges, thereby increasing the chances that a bank will unknowingly be exposed to AI risks by virtue of a relationship with a third party.

To address these issues, banks utilize various tools to detect bias and are investing in red team testing for LLMs. Red teaming provides a credible challenge to the model by seeking vulnerabilities in its conceptualization, design, or application. However, not all banks have the capacity or expertise to deploy these techniques at scale.

In recent years, there has been discussion between bankers and policymakers regarding the potential creation of a public/private standard-setting partnership and corresponding certification program to help reduce the cost, inefficiencies, and uncertainty related to bank onboarding of third-party service providers. A 2020 Request for Information by the Federal Deposit Insurance Corporation requested input as to whether the creation of a public/private partnership could support banks' third-party risk management efforts by certifying or assessing certain aspects of a third-party's products or models or by evaluating a third-party provider's operations or

condition.[44] The FDIC abandoned this work following a change in leadership. However, this concept is worth further exploration, particularly in light of the dramatic evolution of AI.

Increasingly, a bank's ability to compete in the marketplace will depend on its ability to leverage the expertise of third-party service providers—including those that use AI. Banks that are unable to adopt new technologies or partner with new third parties will not be able to provide the products and services that customers expect. Unfortunately, the due diligence necessary to onboard a prospective vendor is costly, inefficient, and time consuming for both banks and service providers. These burdens exist for all institutions and are particularly acute for community banks with limited resources. A standard setting organization could help to address some of these challenges and may encourage third parties to provide increased transparency regarding their use of GAI models.

Some of our member banks also report that select third parties have begun to incorporate their AI policies and information about AI usage into their System of Organization and Controls (SOC) and SOC2 reports. We support this practice and believe that it should be encouraged by policymakers and practitioners. Standardized disclosure of third-party AI practices would help banks to identify, quantify, and manage AI-related risk.

Bank TPRM Programs are Evolving to Account for the Risks Associated with Third Parties That Utilize AI

Banks are concerned about being inadvertently exposed to AI risk stemming from a third-party relationship, and they are in the early stages of taking steps to ensure that third parties are not utilizing AI without the bank's knowledge and consent.

For example, to develop a more fulsome understanding of the extent to which third parties are using AI and the risks that it poses, some banks have expanded, or are in the process of expanding, the AI-related due diligence questions that they submit to *prospective* third parties. They are also requesting that *existing* third parties respond to these AI-focused questionnaires.

In addition to cataloguing and quantifying AI risk, banks are contemplating how AI may impact other aspects of their third-party risk management programs, such as contracting and ongoing third-party oversight. Commonly, contracts with third parties involve multi-year master agreements. The duration of these contracts increases the likelihood of situations where a third party utilizes AI in ways that neither the bank nor the third party envisioned at contract consummation.

Banks report significant difficulties in obtaining contract addendums requiring third parties to provide notice of AI usage. It is unclear the extent to which regulators will lean on banks to seek similar contract modifications for agreements completed prior to the advent of GAI. However, as

[44] https://www.federalregister.gov/documents/2020/07/24/2020-16058/request-for-information-on-standard-setting-and-voluntary-certification-for-models-and-third-party.

a practical matter, this is not a realistic option due to relative bargaining leverage and the sheer numbers of third parties.

In addition, the outsize power of the few Big Tech developers of GAI casts a shadow over negotiations with banks of all sizes. Indeed, the oligopoly of foundational model developers could lead to concentration risk and lack of sufficient resiliency. This system of interdependence could prove calamitous in the event of a system outage or large-scale cyberattack, particularly if AI-enabled services become mission critical (for example, a contact center relying on an internal document management system to respond to customer inquiries). Moreover, AI outputs may be homogenized and lead to missed insights or false assumptions among a critical mass of financial institutions. This vendor concentration of very large foundational LLMs presents a complex situation to address and solve. Swift regulatory action should be tempered with significant industry discussions to strike the right balance between bank requirements and realistic mitigation options. Such coordination is necessary to provide sufficient resilience commensurate with the scale of GAI deployments.

More banks may explore incorporating similar provisions into new contracts, but the ability of banks to negotiate these terms may be limited, particularly when contracting with third parties that serve clients outside of the financial sector that are not subject to the same level of regulatory oversight as banks. If third parties do not agree to proactively communicate their plans to deploy AI, banks anticipate requesting that the third party respond to AI-related oversight questionnaires on a more frequent basis. This could lead to situations where banks learn after the fact that a third party has incorporated AI into its product or service, meaning that the bank's management of AI risk would be reactive rather than proactive.

A third party's utilization of AI in ways that were not envisioned at the time of contracting also has implications for how a bank oversees and monitors that third party. For instance, situations could arise where a third-party's adoption of AI post-contracting requires more extensive model risk and fair lending oversight than the bank anticipated. This scenario would require the bank to devote additional expertise to overseeing the relationship, which skews the cost-benefit analysis that a bank conducts prior to entering a third-party relationship. To some extent, banks encounter this challenge today with respect to monitoring for potential bias associated with AI that has been incorporated into third-party fraud and AML tools. We anticipate that the amount of unplanned monitoring and testing will increase as GAI becomes more commonplace in a bank's third-party relationships.

Banks are also grappling with the potential implications of the risk posed by third parties whose suppliers, subcontractors, or service providers (collectively, Nth parties) are using GAI. The oversight of Nth parties has long been debated in third-party risk management circles, but the issue becomes increasingly complex when GAI is involved. The Interagency Guidance on Third-Party Relationships[45] provides that banks should focus on a third party's own processes for selecting and overseeing its subcontractors and service providers by ensuring that the third party implements effective controls and appropriately manages and mitigates the associated risks. Regulators do not appear to suggest that banks oversee Nth parties directly. However, when AI is

[45] *See* supra, note 11.

involved, a key question is whether the third party is even aware that its suppliers, subcontractors, or service providers are using AI, which would indirectly bring AI considerations into the third party's relationship with the bank. In these situations, it is unlikely that a bank could realistically limit or control its exposure to AI.

Banks will continue to modify their third-party risk management programs as AI evolves. In addition to the challenges discussed above, banks are wrestling with the adequacy of a third party's controls pertaining to AI, the ability to obtain adequate audit rights pertaining to AI, and the ability of the bank to incorporate indemnification provisions into contracts with third parties. The relative market power of the parties will impact a bank's ability to demand these terms, and in any event indemnifications are only as good as the company that provides them. These are long-standing challenges in third-party risk management generally, but they take on increased importance with the advent of GAI.

We cannot overstate the role that third-party risk management plays in ensuring that AI is deployed in a responsible manner. Banks are increasingly likely to be exposed to AI via their third and Nth parties and must have visibility into those AI uses in order to make appropriate business decisions.

E. Illicit Finance

Many banks use AI and machine learning technologies as part of their risk-based approach to BSA and sanctions compliance. In order to adopt an effective risk-based approach, banks must have an accurate understanding of the actual risks associated with their business practices. Banks also need refined and accurate models to avoid expending unnecessary resources investigating false positives; while ensuring they do not miss important red flags. In addition, generative AI can be used by bad actors to create realistic identity documents. Treasury's Financial Crimes Enforcement Network (FinCEN) has warned about the dangers posed by fraudulent identity documents, for example, fraudulent passport cards.[46]

F. Other

In addition to the above, there are several other risks on banks' radars that stem from AI and GAI:

Being left behind- Banks, especially community banks, are concerned with becoming obsolete. They are trying to institute policies, conduct training, and block questionable sites. But how do they build trust with GAI if they block the sites and drive activity underground? A better way is to encourage a legitimate path with proper controls. Monitoring/auditing software can screen for data loss prevention to give notice on how users are employing AI tools they are permissioned. This cannot happen if use occurs by rogue employees. ABA members also addressed the

[46] https://www.fincen.gov/news/news-releases/fincen-issues-notice-use-counterfeit-us-passport-cards-perpetrate-identity-theft.

challenge associated with attracting and retaining data scientists in the current hiring environment. Instead, upskilling traditional business analysts is a more viable path.

Budgeting- A challenge for community banks is how to budget for AI software, since it is now more of a general "cost of doing business." How are funds to be raised and allocated given all the competing business and regulatory priorities?

Digital accessibility- Banks have the obligation to make desktop and mobile platforms accessible from a digital perspective. How do AI and GAI fit in?

Intellectual Property- There is uncertainty of intellectual property rights over AI-generated works and concerns over infringement and rights of use. In addition, there is risk of trade secret information leaks through the use of GAI tools.

Antitrust- As with other technology-neutral requirements, antitrust laws continue to apply to business interactions regardless of whether AI or GAI is employed.

Regulatory- Risks of regulatory ambiguity and/or fragmentation can create challenges for financial institutions and other companies operating in the financial services sector that seek to leverage AI technologies. A lack of clear, consistent governance frameworks – including frameworks that address privacy and security best practices for AI use cases – can create legal risk and higher compliance costs, which can be particularly challenging for smaller institutions and companies to absorb. Regulators must clarify governance and documentation expectations for AI applications that reasonably allows banks of all sizes the flexibility to chart their own path based on risk appetite and business model.

V. Recommendations for Action

Based on the above use cases, opportunities, and risks, ABA and the Associations have several recommendations for policymakers. We highlighted two of these in the introduction[47] but the below goes into significantly more detail and urges additional pursuits.

A. Legislation

We ardently encourage international and multijurisdictional cooperation to pass laws pertaining to AI that are risk-based and industry-focused. Existing law is technology-neutral and applies to use cases whether or not AI is used to deliver them. Lately, state legislatures have been scrambling to pass laws pertaining to the development and use of AI and GAI.

However, Congress must assert its leadership over this issue and pass AI-specific laws that clearly preempt burgeoning state requirements while recognizing banks' unique status as the

[47] First, that any new horizontal federal law pertaining to AI preempt state requirements and clearly exclude banks from any duplicative obligations; second, a call for updated model risk management guidance from the prudential regulators for clarity and to incorporate changes to the ecosystem, but only after an appropriate notice and comment period.

only industry examined for compliance with model risk management expectations. We must avoid the patchwork of state laws that we have witnessed in the comprehensive privacy space given the potential adverse consequences for consumers and national security. Instead, the existing framework applied to banks should be expanded to non-banks and other industries.

To the extent possible, Congress should leverage the work of existing standard setters as it attempts to introduce legislation in the AI realm; for example, to promote consistent nomenclature. This is the only way AI can be meaningfully addressed on a global scale. While we recognize some markets have already advanced AI regulations without global cooperation (e.g., the EU AI Act[48]), we believe regulatory harmonization and flexibility, where possible, will support technology innovation and adoption. Interoperability among the myriad of emerging frameworks is integral. A particular model should not be prioritized simply because it was first on the scene; rather, quality, practicality, and effectiveness should be the criteria.

We support other government activity such as: (1) supporting research activity that would help detect and prevent cyberthreats and fraud; (2) supporting workforce development efforts to ensure the workforce keeps pace with technical advances (e.g., AI-related training and certifications); and (3) strengthening public/private partnerships to increase awareness of cyber and fraud threats.

B. Regulation

Treasury can work with other federal agencies and international financial institutions to help ensure that frameworks governing the use of AI in financial services contexts are interoperable across sectors and jurisdictions. Technology-neutral requirements in civil rights, data privacy and protection, competition, product liability, intellectual property, and numerous other areas of law already apply to AI developers and deployers. Furthermore, payments companies and other entities in the financial sector already operate within a highly complex framework of national and supranational sectoral laws and regulations that can help govern financial-sector AI use cases. It will be critical for Treasury and other federal agencies to (1) leverage existing legal authorities to support responsible, trustworthy AI governance and (2) ensure that any new regulatory activities complement – and do not duplicate or conflict with – existing legal frameworks and obligations. Any new rules resulting from this work should only make clarifications and address any identified gaps to the extent they exist. Further, they should continue to be risk-based and tied to use case.

Regulators should identify clear regulatory outcomes and objectives, while enabling regulated entities the ability to deploy effective risk management techniques based on common standards and best practices. When creating new rules, regulators should consider both current and future use cases with higher inherent risks.

Requirements of banks to examine and monitor third-party AI algorithms, training data, or performance are not possible without third party cooperation. Any rulemaking related to AI should include focus on third-party non-bank AI models, tools, and platforms to impose the same

[48] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689.

obligations directly on such third parties and to require such providers to furnish sufficient credible, reliable information if used by financial institutions or in the financial services market. Regulators should be aware that onerous requirements on banks with respect to using third-party AI may stifle innovation and make it harder for smaller AI developers to compete with larger enterprises that have the resources needed to meet these demands. Accordingly, the agencies should encourage development of voluntary certification programs to provide evidence of compliance for a baseline of fairness, transparency, and explainability. Such programs would aid banks of all sizes and give AI developers incentives to build controls.

At the same time, regulation of AI outside of federally regulated financial institutions needs to be strengthened. If further regulation of AI is provided, it should be designed to level the playing field with under-regulated businesses across the whole AI supply chain, including non-banks and technology companies that provide financial services or support financial services (such as security and IT operations). Regulators should focus on businesses outside the financial industry because they do not have the same prudential regulatory framework as banks and are more likely to create and use AI without guardrails. As has been stated, banks are the only industry with model risk management guidance from regulators.

C. Supervision

ABA and the Associations encourage Treasury to recommend updates to model risk management guidance to be more reflective of bank operations and to make applicability to AI usage more obvious. Existing model risk guidance from the prudential agencies[49] establishes standards for the use of third-party models used by banks, but banks may not have the ability to fully oversee all AI that third parties use to provide services. Updated guidance should clearly delineate the responsibilities when banks use third-party AI and GAI.

Crucially, however, such updates must be developed through a notice and comment period to ensure that the proposed clarifications reflect the current state of technology, industry practice, consumer interests, etc.

Ideally, this updated guidance should address the following areas:

- Revised language should be clearer on the requirements (or lack thereof) for validation of lower risk activities;
- Assessment of the conceptual soundness of GAI models;
- Types and level of performance testing for various LLM types (foundational, purpose built, smaller and trainable open-source, RAG-powered GAI, etc.). Risk treatment may scale depending on the degree of the banks' involvement, and performance validation may require industry benchmark tests as well as specific use case benchmark tests;
- Expectations on the relationship of model risk management with respect to other control domains, such as data and technology risk; and

[49] SR 11-7, OCC Bulletin 2011-12, and FIL-22-2017, supra, note 7.

- Acknowledgement and recognition of the interconnected and complementary nature of the various control domains as part of banks' overall enterprise risk management systems.

The process to assess AI risks may vary depending on the bank and use case, and regulatory risk management expectations should not be one-size-fits-all. Despite the growing variation in AI systems and uses, transparency and fairness remain common requirements in the heavily regulated banking industry. Banks are thus optimal vehicles to apply prudential and ethical AI risk management procedures and assessments to the use of AI – whether developed internally at the bank or by a third party.

Further guidance could help banks differentiate between high- and medium-risk AI use cases and issues and identify appropriate safeguards to use AI responsibly within the existing risk framework. However, overly prescriptive safeguards or prohibitions on the use of AI could become a barrier to valuable use cases if banks do not have appropriate flexibility to test and adapt new AI tools.

D. Other

The financial services industry and its regulators should collaborate to develop standardized strategies for managing AI-related risk. This includes development of standardized disclosure templates for businesses conducting due diligence of third-party usage of GAI; an example might be model cards for validation exercises. Creating sector-specific guidelines based on AI frameworks can lead to more effective mitigation of emerging threats and ensure alignment with regulatory requirements and supervisory expectations. Cooperation between industry and government via public/private partnerships is also needed to meet the challenges posed by advanced technologies.

One example that would help highly regulated institutions and the markets and customers served, is increased transparency and explainability requirements over time, such as independent certifications that the AI model in market use has been appropriately designed and tested, and that potential algorithmic biases have been addressed. However, such certifications should be voluntary and applicable to certain use types and risk levels. Prerequisites for this initiative include uniform definitions of standards and measures, a mechanism for approval of certification bodies, and compliance deference given to such certification.

Federal financial regulators should seek to develop an approach to explainability that transcends traditional methodologies by leveraging a suite of coordinated risk management practices, including but limited to data governance, weighted decision-making criteria, assurance and testing, and continuous risk monitoring. This can be achieved by mapping to the NIST AI RMF (as well as the NIST Privacy Framework[50]) and/or creating sector-specific profiles. This holistic approach should include all participants in the AI ecosystem, including technology companies and non-financial industry actors, particularly because the economics of LLM development prevents internal development thereof and drives adoption of third-party offerings. Such work

[50] https://www.nist.gov/privacy-framework.

could be the foundation of interoperability across sectors and jurisdictions and would allow the entire ecosystem to innovate confidently and responsibly.

Conclusion

ABA, the Associations, and our members (comprised of banks of all sizes) are grateful for the opportunity to provide Treasury with our views on the promise and risks of the expanding use of AI and GAI, including our recommendations for legislative, regulatory, and supervisory actions. We stand ready to work with policymakers on this vital issue.

To that end, ABA and the Associations support Treasury efforts to address AI risks under FSSCC and FBIIC public/private partnerships. In response to the Treasury paper on AI and cyber, these groups are working to develop several workstreams to address a variety of AI-related risks such as identity, authentication, and combating fraud. We look forward to collaborating on other initiatives with Treasury as well as other key stakeholders.

If you have any questions about this comment, please contact Ryan T. Miller (rmiller@aba.com) at (202) 663-7675.

Respectfully submitted,

Ryan T. Miller
Vice President & Senior Counsel, Innovation Policy
American Bankers Association

And on behalf of:

California Bankers Association
Colorado Bankers Association
Delaware Bankers Association
Hawaii Bankers Association
Iowa Bankers Association
Kentucky Bankers Association
Louisiana Bankers Association
Maryland Bankers Association
Massachusetts Bankers Association
Michigan Bankers Association
Nebraska Bankers Association
Nevada Bankers Association
New York Bankers Association
Ohio Bankers League

Pennsylvania Bankers Association
South Dakota Bankers Association
Virginia Bankers Association
Washington Bankers Association
West Virginia Bankers Association
Wisconsin Bankers Association
Wyoming Bankers Association