

For release on delivery
6:15 p.m. EST
November 22, 2024

Artificial Intelligence in the Financial System

Remarks by

Michelle W. Bowman

Member

Board of Governors of the Federal Reserve System

at

The 27th Annual Symposium on Building the Financial System of the 21st Century:
An Agenda for Japan and the United States

Washington, D.C.

November 22, 2024

Discussions of artificial intelligence (AI) inevitably center on two main points: risks and benefits.¹ Both of these can be frustratingly vague and amorphous. Proponents of AI project its widespread adoption will be as momentous as the industrial age—radically improving efficiency, increasing labor productivity, and changing the world economy. Skeptics largely focus on the risks, noting that it may introduce new and unpredictable variables into the economy and the financial system, including new forms of cyber-risk and fraud.

It would be impossible to predict what the future holds for AI, or how its use and impact will evolve over time. But as the technology continues to mature, as new use cases evolve, and it is rolled out more broadly, we will almost certainly be surprised by how it is ultimately used.

Looking at the financial industry-specific implications of AI, it is helpful to consider not only how it may change the financial system, but also how regulatory frameworks should respond to this emerging technology. Are the existing frameworks sufficient? If not, how can regulators best balance the risks AI may pose to bank safety and soundness and financial stability with the need to allow for continued innovation?

Broader availability of generative AI and large language models have created headlines and spiking stock prices, but the financial services sector has been using AI for some time.² Over time, it has become clear that AI's impact could be far-reaching, particularly as the technology becomes more efficient, new sources of data become available, and as AI technology becomes more affordable.

¹ The views expressed here are my own and not necessarily those of my colleagues on the Federal Reserve Board or the Federal Open Market Committee.

² In 2017, the Financial Stability Board was already considering the financial stability implications of AI and machine learning in financial services. See Financial Stability Board, “Artificial intelligence and machine learning in financial services: Market developments and financial stability implications” (Basel: Financial Stability Board, November 2017), <https://www.fsb.org/uploads/P011117.pdf>.

Do We Need a Definition of AI?

Before discussing the implications of AI and regulatory policy approaches, we should ask whether we need a definition of AI. As it has advanced, the number and variety of definitions used to define AI have expanded.³ Some definitions focus on the algorithms—like the use of machines to learn and reason in a way that simulates human intelligence. Others focus on the outputs—the ability to perform complex tasks normally done by humans. In 2021, an interagency request for information from federal banking regulators, including the Federal Reserve, sought comment on banks’ use of AI, but notably avoided using any single definition. Instead, this request listed a few possible use cases, features, and forms. These included the use of structured and unstructured data; the use of alternative data sources; voice recognition and natural language processing; the algorithmic identification of patterns and correlations in training data to generate predictions or categorizations; and “dynamic updating,” where an algorithm has the capacity to update without human intervention.⁴

While each definition of AI may serve its own purpose in the context of how it is used, any single narrow definition can be criticized. A more generic definition runs the risk of oversimplifying the range of activities, use cases, and underlying technology. A definition that captures the variability of AI technology in a more granular way runs the risk of being unwieldy in its length, and obsolete in the short-term as new forms and use cases emerge.

³ The National Artificial Intelligence Initiative Act of 2020, 15 U.S.C. § 9401(3) defines artificial intelligence as “... 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-(A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.”

⁴ See Request for Information and Comment on Financial Institutions’ Use of Artificial Intelligence, Including Machine Learning, 86 Fed. Reg. 16837 (March 31, 2021), <https://www.govinfo.gov/content/pkg/FR-2021-03-31/pdf/2021-06607.pdf>.

Within this definitional question—of whether and how you define AI—lies a more important policy question: Specifically, for what purpose is a definition required? In the context of the financial system, the definition of AI may help to delineate how the regulatory system addresses it and establishes the parameters for how it can be used by regulated institutions. Other specific contexts could also be included, like third-party service providers that support banks or other financial services providers, or use by regulators in support of their mandates.

A definition helps regulators and regulated institutions understand the activities that are subject to rules and requirements by defining the scope. While this definitional question is important to establish clarity about the scope of what constitutes AI, it can also distract us from a more important point—what is the appropriate policy framework to address the introduction and ongoing use of AI in the financial system?

I have no strong feelings about the ideal or optimal definition of AI, and some version of the many definitions floating around are probably adequate for our purposes. At a minimum though, a definition must establish clear parameters about what types of activities and tools are covered. But before leaving the topic, I want to offer a cautionary note. A broad definition of AI arguably captures a wider range of activity and has a longer “lifespan” before it becomes outmoded, and potentially never becomes outdated. But a broad definition also carries the risk of a broad—and undifferentiated—policy response. This vast variability in AI’s uses defies a simple, granular definition, but also suggests that we cannot adopt a one-size-fits-all approach as we consider the future role of AI in the financial system.

Innovation and Competitiveness

Knowing that the technology and use of AI continues to evolve leads to the question of how it should be viewed by regulators, particularly in light of the need for innovation and the effect on competition.

Innovation

AI tools have the potential to substantially enhance the financial industry. In my view, the regulatory system should promote these improvements in a way that is consistent with applicable law and appropriate banking practices.

One of the most common current use cases is in reviewing and summarizing unstructured data. This can include enlisting AI to summarize a single report or to aggregate information from different sources on the same or related topics. The AI “output” in these cases may not directly produce any real-world action, but it provides information in a more usable way to assist a human. AI use cases like this may present opportunities to improve operational efficiency, without introducing substantial new risk into business processes. In some ways, the joining of AI outputs with a human acting as a “filter” or “reality check” can capture efficiency gains and control for some AI risks. Similarly, AI can act as a “filter” or “reality check” on analysis produced by humans, checking for potential errors or biases.

AI tools may also be leveraged to fight fraud. One such use is in combatting check fraud, which has become more prevalent in the banking industry over the last several years. In a recent report, the Financial Crimes Enforcement Network noted that from February to August of 2023, there were over 15,000 reports received related to check fraud, associated with more than \$688

million in transactions (including both actual and attempted fraud).⁵ The growth in check fraud over the past several years has caused significant harm not only to banks and the perceived safety of the banking system but also to consumers who are the victims of fraudulent activity. The regulatory response to help address this growing problem has unfortunately been slow, lacking in coordination, and generally ineffective.

Could AI tools offer a more effective way for banks to fight against this growing fraud trend? We already have some evidence that AI tools are powerful in fighting fraud. The U.S. Treasury Department recently announced that fraud detection tools, including machine learning AI, had resulted in fraud prevention and recovery totaling over \$4 billion in fiscal year 2024, including \$1 billion in recovery related to identification of Treasury check fraud.⁶ While the nature of the fraud may be different in these cases, we should recognize that AI can be a strong anti-fraud tool and provide significant benefits for affected bank customers.

If our regulatory environment is not receptive to the use of AI in these circumstances customers are the ones who suffer. AI will not completely “solve” the problem of fraud—particularly as fraudsters develop more sophisticated ways to exploit this technology. But it could be important if the regulatory framework provides reasonable parameters for its use.

Another often-discussed use case for AI in financial services is in expanding the availability of credit. AI is not the first technology with potential to expand access to credit for the “un-” or “underbanked.” We have long viewed alternative data as a potential opportunity for

⁵ Financial Crimes Enforcement Network, *Financial Trend Analysis: Mail Theft-Related Check Fraud: Threat Pattern & Trend Information, February to August 2023* (Vienna: Financial Crimes Enforcement Network, September 2024), <https://www.fincen.gov/sites/default/files/shared/FTA-Check-Fraud-FINAL508.pdf>.

⁶ See U.S. Department of the Treasury, “Treasury Announces Enhanced Fraud Detection Processes, Including Machine Learning AI, Prevented and Recovered Over \$4 Billion in Fiscal Year 2024,” news release, October 17, 2024, <https://home.treasury.gov/news/press-releases/jy2650>.

some consumers, like those with poor or no credit history but with sufficient cash flow to support loan repayment.⁷

AI could be used to further expand this access, as financial entities mine more data sets and refine their understanding of creditworthiness. Of course, we also know that using AI in this context—in a way that has more direct impact on credit decisions affecting individual customers—also presents more substantial legal compliance challenges than other AI use cases.

AI also has promise to improve public sector operations, including in regulatory agencies. As I have often noted, the data relied on to inform the Federal Open Market Committee decision-making process often is subject to revisions after-the-fact, requiring caution when relying on the data to inform monetary policy.⁸ Perhaps the broader use of AI could act as a check on data reliability, particularly for uncertain or frequently revised economic data, improving the quality of the data that monetary policymakers rely on for decision-making. Additional data as a reliability check or expanded data resources informed by AI could improve the FOMC's monetary policymaking by validating and improving the data on which policymakers rely.

While these use cases present only a subset of the possibilities for the financial system, they illustrate the breadth of potential benefits and risks of adopting an overly cautious approach that chills innovation in the banking system. Over-regulation of AI can itself present risks by

⁷ See Board of Governors of the Federal Reserve System, Consumer Financial Protection Bureau, Federal Deposit Insurance Corporation, National Credit Union Administration, and Office of the Comptroller of the Currency, “Interagency Statement on the Use of Alternative Data in Credit Underwriting,” news release, December 12, 2019, <https://www.federalreserve.gov/supervisionreg/caletters/CA%2019-11%20Letter%20Attachement%20Interagency%20Statement%20on%20the%20Use%20of%20Alternative%20Data%20in%20Credit%20Underwriting.pdf>

⁸ See Michelle W. Bowman, “Perspectives on U.S. Monetary Policy and Bank Capital Reform,” (speech at Policy Exchange, London, England, June 25, 2024) <https://www.federalreserve.gov/newsevents/speech/files/bowman20240625a.pdf>.

preventing the realization of benefits of improved efficiency, lower operational costs, and better fraud prevention and customer service.

Effect on Competition

The regulatory approach and framework can also promote competition in the development and use of AI tools in the financial system.

An overly conservative regulatory approach can skew the competitive landscape by pushing activities outside of the regulated banking system or preventing the use of AI altogether. Inertia often causes regulators to reflexively prefer known practices and existing technology over process change and innovation. The banking sector often suffers from this regulatory skepticism, which can ultimately harm the competitiveness of the U.S. banking sector.

In the United States, we often think about the financial system based on the regulatory “perimeter.” We view institutions within the scope of federal banking regulation (banks and their affiliates) as being “in the perimeter,” while entities that operate under other regulatory frameworks (including money transmitters licensed under state law) are “outside the perimeter.” Of course, the global financial system includes institutions that operate on a cross-border basis, and tools and approaches often permeate throughout the financial system once they have been deployed successfully in some other part of the system. But we know that the regulatory perimeter is permeable, and there is always the risk that activity pushed outside the perimeter can transmit risk back into the system even as the activities garner less scrutiny and regulation than banks. Put differently, the overly conservative approach may present only a façade of safety, masking underlying risks to the financial system and those who rely on it.

Of course, there are risks to being overly permissive in the AI regulatory approach. As with any rapidly evolving technology, supervision of its use should be nimble. Its users must

make sufficient risk-management and compliance investments to conduct activities in a safe and sound manner, and in accordance with applicable laws and regulations. While the banking system has generally been cautious and deliberate in its AI development and rollout, others have not. When left improperly managed and unmonitored, it can result in unintended outcomes and customer harm. For example, certain generative AI models have been known to generate nonsensical or inaccurate outputs, sometimes called “hallucinations.” In some cases, AI hallucinations have not involved significant harm, for example when discovered in a testing environment without customer-facing implications.

The Sufficiency of Existing Regulatory Tools

While helpful to acknowledge the risks of over-regulation and under-regulation, we must understand our currently regulatory stance—what tools do we have to promote AI’s benefits while helping to mitigate the risks? To address this topic, we should widen the lens to consider our approach to innovation broadly. Supporting innovation in the financial system can and should apply to the introduction and use of AI.⁹

When we consider AI risks, many of these are already well-covered by existing frameworks. For example, AI often depends on external parties—cloud computing providers, licensed generative AI technologies, and core service providers—to operate. AI can also pose model risks in the banking context, with associated data management and governance concerns. AI can also impact a bank’s cyber-resiliency as AI-related fraud tools become more widespread and more anti-fraud tools become available.

⁹ See Michelle W. Bowman, “Innovation and the Evolving Financial Landscape” (speech at The Digital Chamber DC Blockchain Summit 2024, Washington, D.C., May 15, 2024), <https://www.federalreserve.gov/newsevents/speech/files/bowman20240515a.pdf>.

While AI may be on the frontier of technology, it does not operate outside the existing legal and regulatory framework. AI is not exempt from current legal and regulatory requirements, nor is its use exempt from scrutiny. This is particularly true with AI; its use must comply with current laws and regulations, including fair lending, cyber security, data privacy, third-party risk management, and copyright. And when AI is deployed in a bank, an even broader set of requirements may apply depending on the use case.

Regulators are often playing “catch-up” with banks at the forefront of innovation. As a result, they often suffer from significant disadvantages in terms of understanding how the technology works, understanding the uses of AI within financial institutions, and keeping up-to-date with the latest AI developments. Of course, further compounding this challenge is that much of the work in AI innovation occurs far outside the banking system, including in the development and testing of generative AI models and in compiling the data sources on which to train the models.

Despite these challenges—and the understandable regulatory instinct to limit its use in the financial system—we must avoid this temptation. A few general principles should govern a coherent regulatory approach, which are the same principles that I apply to innovation generally.¹⁰

First, we must understand AI before we consider whether and how to change our regulatory approach. With respect to various internal use cases, the Board has published a compliance program that governs artificial intelligence.¹¹ One of the foundational elements for a

¹⁰ See Bowman, “Innovation and the Evolving Financial Landscape.”

¹¹ See Board of Governors of the Federal Reserve System, *Compliance Plan for OMB Memorandum M-24-10* (Washington: Board of Governors, September 2024), <https://www.federalreserve.gov/publications/files/compliance-plan-for-omb-memorandum-m-24-10-202409.pdf>.

successful approach to AI, and one mentioned in this plan, is the development and acquisition of staff expertise.

Many banks have increased AI adoption to an expanding number of use cases. As this technology becomes more widely adopted throughout the financial system, it is critical that we have a coherent and rational policy approach. That starts with our ability to understand the technology, including both the algorithms underlying its use and the possible implications—both good and bad—for banks and their customers.

In suggesting that we grow our understanding and staff expertise as a baseline, I acknowledge that this has been, and is likely to remain, a challenge. The Federal Reserve and other banking regulators compete for the same limited pool of talent as private industry. But we must prioritize improving our understanding and capacity as this technology continues to become more widely adopted.

Second, we must have an openness to the adoption of AI. We need to have a receptivity to the use of this technology and know that successful adoption requires communication and transparency between regulated firms and regulators. One approach regulators can use to reframe questions around AI (and innovation generally) is to adopt a posture that I think of as technology agnosticism.

We should avoid fixating on the technology, and instead focus on the risks presented by different use cases. These risks may be influenced by a number of factors, including the scope and consequences of the use case, the underlying data relied on, and the capability of a firm to appropriately manage these risks. Putting activities together may be a helpful way to get a sense of broad trends (for example, the speed of AI adoption in the industry), but is inefficient as a way to address regulatory concerns (like safety and soundness, and financial stability). This may

seem like an obvious point, but at times regulators have fallen prey to overbroad categorizations, treating a diverse set of activities as uniformly and equally risky.

This approach allows us to be risk-focused, much like we try to do with other forms of supervision, moderating intensity for low-risk activities, and increasing the intensity for higher-risk ones.

Of course, regulatory agencies do not operate in a vacuum, so we must also ask what type of coordination we need to ensure that we promote safe and sound adoption of AI, and address broader financial stability risks, both domestically and internationally. As a threshold matter, we need coordination both within each agency and among domestic regulators that play a role in the supervision and regulation of the financial system, which requires an environment of open sharing of information.

A posture of openness to AI requires caution when adding to the body of regulation. Specifically, I think we need a gap analysis to determine if there are regulatory gaps or blind spots that could require additional regulation and whether the current framework is fit for purpose. Fundamentally though, the variability in the technology will almost certainly require a degree of flexibility in regulatory approach.

Closing Thoughts

Before closing, I want to thank the organizers of this event for the invitation to address you this evening, and to thank the many speakers and participants who have contributed to the symposium.

Artificial intelligence has tremendous potential to reshape the financial services industry and the broader world economy. While I have suggested in my remarks that we need not rush to regulate, it is important that we continue to monitor developments in AI and their real-world

effects. In the long run, AI has the potential to impact many aspects of the Fed's work, from our role in supervising the payment system, to the important work we do promoting the safe and sound operation of banks and financial stability. AI may also play a growing role in monetary policy discussions, as the introduction of AI tools alter labor markets, affecting productivity and potentially the natural rate of unemployment and the natural rate of interest.

But as we engage in ongoing monitoring—and expand our understanding of AI technology and how it fits within the bank regulatory framework—I think it is important to preserve the ability of banks to innovate and allow the banking system to realize the benefits of this new technology.