For release on delivery
5:00 a.m. EST (11:00 a.m. local time)
November 16, 2026

Some Potential Financial Stability Implications of Artificial Intelligence

Remarks by

Michael S. Barr

Member

Board of Governors of the Federal Reserve System

at the

Seventh Annual Association for Computing Machinery
International Conference on Artificial Intelligence in Finance

Milan, Italy

November 16, 2026

Thank you for the opportunity to speak to you today.[1] This is the fourth speech I have given in the past year devoted exclusively to artificial intelligence (AI), and most of the others have also addressed it in some way. The effects of AI on people and economies around the globe are among the most important issues we are facing today.

Let me stipulate at the outset that I am an optimist about the long-term promise of AI to empower individuals, businesses, and institutions across society. I remain hopeful that if we manage today's challenges and tomorrow's potential disruptions forcefully and effectively, that AI will lift up people, make them more productive, and broadly raise living standards—that most people will learn more *and* earn more. One of the reasons I'm optimistic about this technology is that I'm also optimistic about the capacity of society to understand and respond to the changes that AI will bring—to engage in a constructive dialogue about ensuring that the benefits from AI are not overwhelmed by the disruptions it will inevitably bring and the risks it may pose.

This conference is part of that dialogue, directed at questions that I believe will be crucial in ensuring that AI is a positive force in the financial system that empowers individuals to better manage their finances and access credit and services, lowers costs and enhances the capabilities of financial service providers, and also makes markets function more effectively. But AI can only succeed in strengthening our financial system in this way if it also supports financial stability, something that is indispensable to a healthy economy.

As has been the case for some time, there are important insights about AI and financial stability in the Federal Reserve's latest *Financial Stability Report* (FSR) issued

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

on Friday. This edition of the twice-yearly report focuses on the AI investment boom and research by the Fed's staff to assess the degree of leverage in some of the innovative debt structures being employed to finance this growth. I will recap the report's findings and then turn to a discussion of what economic research can tell us about the benefits and risks of integrating AI into financial markets, focusing on wholesale trading. Next, I will examine the potential benefits and risks of AI's effect on cybersecurity and then finally consider what AI might mean for consumer and investor protection.

The AI Investment Boom

Anyone watching the news or their retirement plans these past couple of weeks is aware of the questions that investors have raised about the sustainability of the AI investment boom. I should issue the usual disclaimer here that policymakers on the Federal Reserve Board do not target levels of financial markets. Markets go up, and markets go down, and I assiduously avoid commenting on the ups and downs in markets. We do assess longer-term issues that relate to financial stability, and the latest FSR covers several.

In a previous speech, I discussed three different scenarios summarizing the wide range of possible outcomes for how much AI could alter the economy.[2] In one possible scenario among several, even if AI turns out to be a transformative innovation that vastly improves productivity, a key question is whether investors will see returns on the AI buildout on a timeline and to an extent consistent with their expectations or whether a different outcome could lead to a significant repricing. The potential investment losses

[2] See Michael S. Barr (2026), "Will Artificial Intelligence Broadly Raise Living Standards or Drive Income and Wealth Inequality?" speech delivered at "Next-Gen Financial Inclusion," the third annual Financial Inclusion Conference hosted by the Federal Reserve Board, Washington, July 14, https://www.federalreserve.gov/newsevents/speech/barr20260714a.htm.

that could result from a large-scale repricing could lower output growth in two ways: first, from the direct effect of a drop in investment, which is a component of gross domestic product; and, second, from the knock-on effects of this loss of wealth in lowering consumer and potentially business spending. As discussed in the latest FSR, the equity premium, a rough measure of the additional return that investors require for holding stocks relative to risk-free bonds, is at low levels relative to historical norms, though not nearly as low as the levels reached during the dot-com stock market boom.[3] This is just one indicator to assess where broad measures of equity valuations are relative to historical norms.

Leverage in the Data Center Buildout

In addition to the question of equity valuations, the AI buildout has been increasingly funded by debt. The Federal Reserve's FSR delves into three ways in which the financing of AI data centers and related investments, which it calls "advanced computing debt," could affect financial stability: (1) the rapid growth of this debt, (2) circular financing networks, and (3) other complex financial obligations. In the case of a significant repricing of AI investments, stress from these three aspects of AI financing could amplify the repricing of AI investments.

On-balance-sheet debt of the seven largest AI firms has been growing significantly in 2026 and is expected to exceed $800 billion before year's end. The growth of debt at these firms brings to mind the growth of debt in the telecom investment cycle of the late 1990s and early 2000s. The FSR compares hyperscaler firms'

[3] See figure 1.5 in Board of Governors of the Federal Reserve System (2026), *Financial Stability Report* (Washington: Board of Governors, November), p. 8, https://www.federalreserve.gov/publications/files/financial-stability-report-20261113.pdf

cumulative capital spending and debt issuance since 2020 with those of all firms in the telecommunications industry during the broadband investment cycle that started in 1995. While the hyperscalers have yet to surpass their comparators of the earlier boom, the recent acceleration in their capital expenditures suggests they may do so soon.[4]

While these firms' on-balance-sheet debt levels have grown substantially, their other obligations, such as long-term contracts, lease commitments, and residual value guarantees, have grown even more rapidly. The FSR estimates that the potential maximum nominal value of these types of obligations, or off-balance-sheet liabilities, now exceeds $3 trillion.[5] A significant share of these obligations results from complex arrangements in which hyperscalers and other large firms with strong credit ratings indirectly support the financing of projects that involve other firms with weaker credit, or provide support for public utility investments. These off-balance-sheet liabilities include obligations relating to energy supply contracts, long-term data center and compute leases, and a wide variety of off-balance-sheet financing structures. As I have discussed in previous speeches, this financial complexity can be opaque to unsophisticated investors, and that opacity can exacerbate selloffs when investor sentiment shifts, as it is hard to know where risk resides. As a result, the development of these substantial off-balance-sheet structures warrants close monitoring. Perhaps just as important, in addition to the hyperscalers, lower-credit-quality firms most dependent on external funding have seen

[4] See the box "Leverage in the Data Center Buildout" in Board of Governors of the Federal Reserve System (2026), *Financial Stability Report* (Washington: Board of Governors, November), pp. 23-25, https://www.federalreserve.gov/publications/files/financial-stability-report-20261113.pdf

[5] Note that firms generally present off-balance-sheet liabilities before discounting, which complicates comparisons to on-balance-sheet numbers, which are discounted.

their cost of capital rise significantly in recent months and, as a result, are turning to more complex asset-backed financings.

In addition, the hyperscaler firms and large graphics processing unit (GPU) producers also have over $100 billion in cumulative equity investments in large language model (LLM) labs and cloud-computing start-ups. These equity investments represent a form of "circular financing" in which hyperscaler firms and GPU makers invest capital in smaller firms, and the smaller firms turn around and use these funds to rent or purchase GPUs and compute from the firm that made the investment. These investments represent a relatively small share of the hyperscalers' and GPU producers' balance sheets, but they are critical for the smaller firms receiving the investments. For the AI buildout investments to ultimately become economically viable, the smaller firms will eventually have to replace that funding from their suppliers, the hyperscalers and GPU makers, with revenue from paying customers.

Another indirect form of circular financing in which the hyperscalers and GPU producers support purchases of their products and services is the use of collateral guarantees for asset-based financings. As the firms most dependent on external financing bump up against financing constraints, they shift from issuing corporate bonds to issuing asset-based financing. To support the credit of these financings, the larger higher-credit firms often guarantee the resale value for the data centers or GPUs that are used as collateral to back the financings. These guarantees, or collateral backstops, are reported as off-balance-sheet liabilities because they represent the degree to which hyperscalers and GPU producers might absorb losses in the event of defaults and collateral liquidation.

As we take a step back and consider the enormous scope of AI-related investment in capital stock, it is helpful to consider how it compares to that for previous technology cycles. By one estimate, the capital investment associated with the AI buildout could plausibly rival or even exceed the investments associated with previous infrastructure investment booms, including those associated with railroads in the 1870s, highways in the 1950s and '60s, and fiber-optic cable in the 1990s.[6] Each of these investment booms may have felt overbuilt at some point, but today we are glad to have the railroad, highway, and fiber-optic capacity that they put in place. Some of those earlier booms fizzled, with consequences for the broader financial system. Given the still-limited amount of debt backing the current boom, I'm less concerned about direct losses. However, with leverage at hedge funds and broader risk appetite apparently quite high, a sharp reversal could lead to rapid deleveraging and, potentially, self-reinforcing fire sales. Further, the leverage used in equity strategies has also increased in recent years. This leverage is provided, in part, via the prime brokerage services of large dealers. Thus, while the high levels of capital at the largest banks provide a bulwark against spillovers, counterparty exposures could rise in a period of equity market volatility, so it's important to keep monitoring exposures, including exposures to counterparties that have taken large positions in AI.

Moreover, the bulk of a data center's cost comes from the GPUs and other specialized chips used in AI. One concern is that these specialized chips may depreciate quickly as newer technologies overtake them, such that the capacity they represent will

[6] See Stijn Van Nieuwerburgh (2026), "Financing the AI Buildout," Brookings Papers on Economic Activity, Fall, pp. 1–28, https://www.brookings.edu/wp-content/uploads/2026/09/4c_Van-Nieuwerburgh.pdf.

not be of much use in the future. Currently, the lease rates of older vintages of chips appear to be holding up, suggesting that there are uses even for computing technology behind the cutting edge, but those trends are worth monitoring. It is hard to say what the other parts of the AI-related capital stock—substations, high-voltage connections, and grid interconnection—may be used for if demand for their services disappoints. In some areas of the country, no doubt, extra power and network capacity will be useful regardless of the future of AI. In other areas, it may take years for excess capacity to become useful, as was the case with "dark fiber" in some areas after the 1990s boom and bust. Or maybe there will be no AI investment fizzle but a continued boom, and our demand for AI services will make full use of the capital being installed now and then some. Only time will tell.

Unintended Consequences of AI in Wholesale Securities Trading

Another possible effect of AI on the financial system is the technology's potential to reshape trading in financial markets, a topic addressed by previous FSRs.[7] Generative AI is a powerful tool that can process broader and richer data and information and apply more complex logic than traditional trading algorithms. It can lower barriers to accessing advanced trading technology and increase diversity of market participants. And it can also help improve market surveillance. Each of these capacities can improve market efficiency, integrity, and functioning.

[7] See the box "Artificial Intelligence and Algorithmic Trading" in Board of Governors of the Federal Reserve System (2025), *Financial Stability Report* (Washington: Board of Governors, November), pp. 15–16, https://www.federalreserve.gov/publications/files/financial-stability-report-20251107.pdf. These issues have also been highlighted by Governor Cook; see Lisa D. Cook (2025), "A Policymaker's View of Financial Stability," speech delivered at Georgetown University's McDonough School of Business Psaros Center for Financial Markets and Policy, Washington, November 20, https://www.federalreserve.gov/newsevents/speech/cook20251120a.htm.

As with any new technology, however, there are risks. While it is my understanding that generative AI is not yet being used for autonomous trading at large scale, we need to be prepared to understand the risks and implications of its adoption.[8] Many of the most important questions about how AI will affect financial markets will be answered not in the abstract, but in the technical choices made about model design, training data, testing, and deployment.

While AI can democratize access to trading expertise and technology, it could also increase information asymmetry and market concentration.[9] The state-of-the art systems are expensive and require massive investment.[10] This could confer an advantage on the largest firms and concentrate market power among them.[11] Even when AI tools are widely available, effectively using them requires considerable skill, potentially favoring larger incumbents.[12]

---

[8] See Board of Governors of the Federal Reserve System (2025), *Financial Stability Report* (Washington: Board of Governors, November), https://www.federalreserve.gov/publications/files/financial-stability-report-20251107.pdf; and "Advances in Artificial Intelligence: Implications for Capital Market Activities," in International Monetary Fund (2024), *Global Financial Stability Report: Steadying the Course: Uncertainty, Artificial Intelligence, and Financial Stability* (Washington: IMF, October), pp. 77–105, https://www.imf.org/en/Publications/GFSR/Issues/2024/10/22/global-financial-stability-report-october-2024.

[9] One study this year showed that ChatGPT helps reduce the information gap faced by retail investors by enabling them to better process publicly available information; see Anne Yanru Chang, Xi Dong, Xiumin Martin, and Changyun Zhou (2026), "AI Democratization and Trading Inequality," *Journal of Accounting Research*, vol. 64 (June), pp. 1287–331.

It is well known that trading across many markets is dominated by a relatively small number of large firms; see John Coughlan and Alexei G. Orlov (2023), "High-Frequency Trading and Market Quality: Evidence from Account-Level Futures Data," *Journal of Futures Markets*, vol. 43 (August), pp. 1126–60; and International Monetary Fund, "Advances in Artificial Intelligence," pp. 77–105 (see note 8).

[10] See Katherine Doherty, Dina Bass, and Justina Lee (2026), "Jane Street and Rivals Are Now Fighting for AI Power as Well as Pace," *Bloomberg,* October 22.

[11] See Thierry Foucault, Leonardo Gambacorta, Wei Jiang, and Xavier Vives (2025), *Artificial Intelligence in Finance: The Future of Banking 7,* Centre for Economic Policy Research (Paris and London: CEPR), https://cepr.org/publications/books-and-reports/barcelona-7-artificial-intelligence-finance.

[12] See Jinfei Sheng, Zheng Sun, Baozhong Yang, and Alan L. Zhang (2026), "Generative AI and Asset Management," *Review of Financial Studies*, pp. 1–48.

The implications for price efficiency of the integration of AI into financial market trading are ambiguous.[13] In his 2026 Jackson Hole paper, Markus Brunnermeier cautioned that securities prices may also become less informative because of asymmetric understanding: AI agents understand humans and can anticipate their actions, but humans do not understand AI agents as well and cannot anticipate what they might do.[14]

Market Herding and a Possible Stampede

If traders rely on a small set of similar AI models trained on similar data, they face the risk of herding behavior—trading strategies could become similar and synchronize reactions to news and other information.[15] Algorithm flaws can cause failures, but a market with large numbers of similar models that may have the same flaw risks failures with systemic consequences.[16] Moreover, bad actors could intentionally

---

[13] AI can improve the information content of prices by incorporating a broader set of information more quickly; see Qiang Cheng, Pengkai Lin, and Yue Zhao (2025), "Does Generative AI Facilitate Investor Trading? Early Evidence from ChatGPT Outages," *Journal of Accounting and Economics*, vol. 80 (November–December), 101821; and Foucault and others, *Artificial Intelligence in Finance* (see note 11). But there is also the risk that AI-generated content on finance-related social media could help spread misleading information, leading to higher volatility; see David Hirshleifer, Lin Peng, Qiguang Wang, Weichen Zhang, and Xiaoyan Zhang (2026), "AI, Opinion Ecosystems, and Finance," NBER Working Paper Series 34807 (Cambridge, Mass.: National Bureau of Economic Research, February), https://www.nber.org/papers/w34807.

[14] See Markus K. Brunnermeier (2026), "Artificial Intelligence and the Brave New World in Finance," paper prepared for the Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyo., August 27–29, https://www.kansascityfed.org/documents/18595/brunnermeier.pdf.

Progress is being made in understanding generative AI models, with research moving from merely observing their internal workings toward experimentally identifying some of the mechanisms that drive their behavior; see Or Shafran, Atticus Geiger, and Mor Geva (2026), "Constructing Interpretable Features from Compositional Neuron Groups," proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 42326–48, https://aclanthology.org/2026.acl-long.1959.pdf.

[15] See Stephen G. Cecchetti, Robin L. Lumsdaine, Tuomas Peltonen, and Antonio Sánchez Serrano (2025), *Artificial Intelligence and Systemic Risk,* Advisory Scientific Committee Reports 2025/16 (Frankfurt, Germany: European Systemic Risk Board, December), http://dx.doi.org/10.2139/ssrn.5866223; and Foucault and others, *Artificial Intelligence in Finance* (see note 11).

[16] See Cecchetti and others, *Artificial Intelligence and Systemic Risk* (see note 15).

manipulate trades, quotes, or news and potentially mislead similar AI trading systems at the same time, amplifying market disruptions.[17]

AI models may echo unfortunate human tendencies. For example, some theoretical research shows that some reinforcement-learning algorithms, so called Q-learning trading algorithms, can autonomously develop collusive behavior without explicit intent.[18] Such self-learning algorithms can also autonomously learn to manipulate markets—for example, through spoofing, which is placing orders designed to create a false impression of market demand—or adopt trading strategies that create and exploit bubbles, thereby destabilizing markets.[19] To attempt to combat such potential risks, exchanges are deploying sophisticated machine learning and AI tools themselves

[17] See Xiang-Li Lim, Puja Singh, and Richard Stobo (2025), *Regulatory Considerations Regarding Accelerated Use of AI in Securities Markets*, IMF Technical Notes and Manuals 2025/016 (Washington: International Monetary Fund), https://www.imf.org/en/publications/tnm/issues/2025/12/24/regulatory-considerations-regarding-accelerated-use-of-ai-in-securities-markets-571909.

[18] See Winston Wei Dou, Itay Goldstein, and Yan Ji (2025), "AI-Powered Trading, Algorithmic Collusion, and Price Efficiency," NBER Working Paper 34054 (Cambridge, Mass.: National Bureau of Economic Research, July), https://www.nber.org/papers/w34054; and Jean-Edouard Colliard, Thierry Foucault, and Stefano Lovo (2026), "Algorithmic Pricing and Liquidity in Securities Markets," *Review of Financial Studies*, pp. 1–49. The likelihood of collusion, however, is smaller when algorithms' learning processes differ. For example, Q-learning algorithms can be more prone to coordination than LLM-based algorithms; see Kartik Anand, Sophia Kazinnik, Agnese Leonello, and Ettore Panetti (2026), "Ex Machina: Financial Stability in the Age of Artificial Intelligence," ECB Working Paper Series 3225 (Frankfurt, Germany: European Central Bank), https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp3225~52d8adae99.en.pdf. Moreover, algorithmic traders have strong incentives to differentiate their strategies, as not colluding can be highly profitable when others collude. This might favor algorithmic heterogeneity as an equilibrium outcome; see Laura Veldkamp (2024), "Discussion of 'AI-Powered Trading, Algorithmic Collusion, and Price Efficiency,' " NBER Summer Institute 2024: Asset Pricing (Cambridge, Mass.: National Bureau of Economic Research, July 11–12), https://conference.nber.org/conf_papers/f201676.pdf.

[19] See Álvaro Cartea, Patrick Chang, and Gabriel García-Arenas (2025), "Spoofing and Manipulating Order Books with Learning Algorithms," working paper, http://dx.doi.org/10.2139/ssrn.4639959; and Winston Wei Dou, Itay Goldstein, and Yan Ji (2026), "Financial Market Fragility in the Era of AI Planning," Wharton School Research Paper, September 16, http://dx.doi.org/10.2139/ssrn.5763222.

for market surveillance.[20] Some trading venues are already acting to reduce risks from the "black box" nature of AI-enabled trading algorithms.[21]

Since we all have an interest in safe, fair, and efficient markets, we need to think through the risks and implications before any new technology is integrated into trading activity or market infrastructure more broadly. When that technology is still not fully understood, and, in fact, its development may be galloping ahead of our understanding, some caution is warranted.

AI and Cyber Risk

Incidents this year have shown the risks from AI models acting autonomously, not in alignment with the goals of their human creators, including cyber risks. Given a goal, AI models have treated security boundaries as obstacles to overcome rather than rules to be followed, raising questions about how effective rules will be even against AI agents instructed to follow them. AI models have exploited some of the same vulnerabilities found by human cyberattackers, such as the heavy reliance of many competing firms on a shared platform provider. AI-enabled cyberattackers have demonstrated that they are able to exploit vulnerabilities that were previously thought too difficult to overcome, and thus not worth fixing, even if they were known in advance. And AI models are also revealing vulnerabilities no one dreamed of just a few years ago. I have spoken before, for example, about how AI can raise cyber risks through its use in creating deepfakes.[22]

---

[20] See Pedro Gurrola-Pérez and Kaitao Lin (2024), "An Analysis of Market Manipulation Definitions Around the World," working paper (London: World Federation of Exchanges, June).
[21] See the CME Group's "Market Regulation Advisory Notice" at https://www.cmegroup.com/content/dam/cmegroup/notices/market-regulation/2024/07/CME-Group-RA2403-5.pdf.
[22] See Michael S. Barr (2025), "Deepfakes and the AI Arms Race in Bank Cybersecurity," speech delivered at the Federal Reserve Bank of New York, New York, April 17, https://www.federalreserve.gov/newsevents/speech/barr20250417a.htm.

Just like the attackers, cyber defenders are employing AI, finding and patching vulnerabilities much faster than humans ever could. But recent incidents raise the question of whether defenders could always be at a disadvantage. Attackers only need to be right once, while defenders need to be right all the time. Cyber defense must be consistently successful, while cyber offense can experiment and fail repeatedly until something works.

All of this is to say that the work of defenders is hard and getting harder given the speed of model innovation. The ability of both the private sector and government to manage cyber risk and effectively manage events as they occur will remain vital to maintaining financial stability. We all have a stake in winning this race against cyberattackers, because none of us can succeed without a strong and safe financial system. I hope those competing for a piece of the future of AI in finance can collaborate to help cyber defense stay ahead of cyberattacks.

AI in Consumer Finance

AI is likely to play a growing role in consumer finance, as it will in most retail businesses, especially those where interactions with customers don't need to involve a human. AI got a head start in finance over many other sectors in providing customer assistance, and consumers are themselves driving adoption, employing chatbots, for example, as a source of investment advice. In many cases, this advice is being dispensed without the structured protections, regulation, and oversight applied when the source of that advice is a financial institution. In addition to individuals turning to chatbots on their own, financial-sector firms are using AI to improve their financial advice.

The increasing use of AI by customers and financial firms raises important investor-protection concerns, such as conflicts of interest, the quality of investment recommendations, and firms' obligations to act in investors' best interests. AI systems may generate recommendations based not only on an investor's financial circumstances and investment objectives, but also on factors that advance a firm's commercial interests, such as promoting proprietary products, generating transaction fees, or encouraging more frequent trading. These conflicts may be embedded in algorithms and difficult for brokers and advisers, investors, or regulators to identify or understand. In addition, AI systems may rely on incomplete or biased data, generate inaccurate information, or recommend investments that are inconsistent with an investor's risk tolerance, financial circumstances, or investment horizon.

Institutions may increasingly use AI to expand household credit access through alternative data analysis. However, AI poses significant compliance challenges. The technology can function as a black box, making it difficult to explain adverse decisions or ensure compliance with fair lending laws.[23] Alternative data sources such as rental payments, utility bills, or employment patterns may expand access to credit if they are used appropriately, but they may also inadvertently serve as proxies for protected characteristics like race or national origin. Thus, care must be taken. Unlike traditional credit reports, consumers often cannot review or dispute this information.

To use AI responsibly, institutions must invest in explainable AI techniques, implement continuous fairness monitoring, maintain detailed model validation records, and provide accessible dispute mechanisms. AI providers must identify biases in training

[23] Such as the Fair Housing Act and the Equal Opportunity Act.

data and monitor outputs to prevent amplifying inequalities or mispricing risks. Those of us who enforce consumer protection laws must remain equally vigilant.

New risks continue to surface. For example, major banks now report finding AI agents attempting to take actions in retail customer accounts, without the bank being able to verify authenticity, consent, and alignment with customer instructions. These incidents underscore that AI in financial services can only succeed long term with appropriate guardrails: strict authentication, explicit consent verification for specific actions, technical controls limiting agent capabilities, and both rules- and values-based constraints. We need technological innovations that ensure AI agents follow established rules and values, especially critical as agentic commerce expands. In addition, AI agents may accelerate deposit outflows during stress periods. Where customers once needed to execute transfers, AI agents could instantaneously detect yield differentials or stability signals and autonomously move funds at high speed, potentially accelerating runs that unfold in minutes, overwhelming traditional liquidity management and regulatory intervention timelines.

Conclusion

I have discussed a number of issues related to integrating AI into the financial system. Among the key challenges are the practical constraints of using AI for core business activities in the financial sector, as firms need to ensure that the resulting processes and outcomes are consistent with relevant laws and appropriate risk management. Large institutions are exploring the use of generative AI, including agentic AI, in their business models—but doing so requires care. To successfully leverage the potential of generative AI on a sustainable basis, decisions based on those models must

be well controlled, numerically and legally precise, explainable, and replicable. AI developers still struggle to some extent with all of those criteria. We need to reduce the risk that AI reinforces biases in consumer lending. And we also need to guard against the risks that could result from the use of AI in financial markets including collusion, market manipulation, or trading strategies that result in significant market volatility or even systemic risk. Opportunities abound for using AI to improve access and efficiency in our financial system in ways that benefit providers and customers. I believe we can and must ensure that this proceeds in ways that protect users and support financial stability.