

For release on delivery
3:55 p.m. EDT (12:55 p.m. PDT)
May 27, 2026

The Opportunities and Risks AI Presents for the Economy and Financial System

Remarks by

Lisa D. Cook

Member

Board of Governors of the Federal Reserve System

at

Stanford Institute for Economic Policy Research

Stanford University

Stanford, California

May 27, 2026

Thank you, Neale, for that kind introduction. Being back on Stanford's campus is always an honor and conjures up great memories. I spent several formative years here—first as a student in the AEA Summer Program, which prepares students to pursue graduate study in economics, and then as a National Fellow at the Hoover Institution. To say that these stints at Stanford were transformative would be an understatement. The summer program prepared me for and set me on a new intellectual and career journey, and my three years here as a postdoc set out an entirely new line of research inquiry. In fact, I started my research on patents and innovation or the economics of innovation here and benefitted greatly from my interaction with economists here at the Stanford Institute for Economic Policy Research (SIEPR), the economics department, the business school, the law school, and Hoover, including Kenneth Arrow, Tim Bresnahan, Jeremy Bulow, Milton Friedman, Avner Grief, Mitch Polinsky, Paul Romer, and Gavin Wright. From my decade spent in the Bay Area—here and at Berkeley, I witnessed how seriously new ideas are taken, examined, implemented, and spread. It is always invigorating to return to such a center of innovation.¹

I applaud SIEPR for holding this event to discuss artificial intelligence (AI) and its power to influence the trajectory of the economy and transform the financial system. I know many in this room are grappling with how to harness this technology's obvious multidimensional promise while being mindful of important risks. Having adopted machine learning in the AEA Summer Program in 2018 when I was director and having used it in my research before coming to the Fed, I arrived at the Board of Governors in

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.

2022 raising questions about and urging the study and adoption of AI. So, rest assured, policymakers at the Federal Reserve are also deeply engaged.

Today, I will start by offering my latest economic outlook, with a focus on implications of AI for both sides of our dual mandate of maximum employment and price stability. Then, consistent with my long-standing support for responsible innovation, I will address the benefits AI could deliver for the financial system before addressing some of the risks and vulnerabilities the technology presents to financial stability. I will conclude by sharing how the Fed itself is embracing the power of AI to help ensure the financial system remains sound and resilient.

Economic Outlook

To set the stage, I will begin with my economic outlook.

Allow me to begin with inflation. Inflation is clearly moving in the wrong direction. Based on the latest data, it is estimated that the personal consumption expenditures (PCE) price index rose 3.8 percent over the 12 months ending in April. That reading is well above our 2 percent target. The recent rise in gasoline prices due to the conflict in Iran was the primary driver. But even when one excludes the volatile food and energy components, core PCE inflation is estimated to have risen by 3.3 percent over the 12 months ending in April—its highest reading since 2023.

Inflation has been pushed up by shocks that should, in theory, be temporary and short lived. Tariffs should result in only a one-time shift upwards in the price level, and the effects of tariffs on inflation should begin to abate soon. The path of energy prices is tied to the ongoing conflict, the results of which are highly uncertain. Still, most

forecasters and market participants, as reflected in oil futures, expect that oil and gasoline prices should decline, to some extent, by the end of the year.

Nonetheless, even temporary and short-lived shocks could influence inflation over the medium term. Firms may embed these shocks into their pricing decisions, and workers may incorporate them into wage negotiations. Moreover, yet another shock to prices could be layered on from the heightened investment demand due to AI. To date, companies have announced more than \$1.5 trillion in data-center plans, only a small portion of which have been realized.² Those figures suggest that substantial AI-related investment remains in the pipeline from data centers alone. Effects of this demand on prices are apparent. Prices have risen significantly for chips, other high-tech equipment, and software. Wages in specialty trades in construction have picked up notably. Electricity and water prices have each increased by about 5 percent over the past year. Further, in the coming years, firms may expand along the intensive margin, but they may also expand along the extensive margin and undertake new AI-related capital expenditure, such as in robotics.

In contrast to inflation, the labor market appears to be largely stable. The unemployment rate, 4.3 percent in April, has remained unchanged, on net, since last summer. The rate is in line with estimates of the natural rate of unemployment, suggesting that the supply and demand of labor are roughly balanced. Despite some high-profile announcements of layoffs, initial claims for unemployment insurance remain

² See Eirik Eylands Brandsaas, Daniel Garcia, Robert Kurtzman, Joseph Nichols, and Adelia Zytek (2025), “Estimating Aggregate Data Center Investment with Project-Level Data,” Finance and Economics Discussion Series 2025-109 (Washington: Board of Governors of the Federal Reserve System, December), <https://doi.org/10.17016/FEDS.2025.109>. For updated data and publicly available results, see Eirik Brandsaas (2026), “Estimating Aggregate Data Center Investment with Project-Level Data,” DataCenterPublic, GitHub repository, <https://github.com/eirikbrandsaas/DataCenterPublic>.

low and stable. However, I view the downside risks to the labor market as being elevated. One factor is heightened uncertainty about output due to the Middle East conflict. A softening in demand could lead to a softening in the labor market. Uncertainty may also weigh on firms' hiring plans, which may be one reason for the current low-hire environment.

Furthermore, I have been and will continue to be highly attentive to AI developments and how they will affect the labor market. We could be approaching the most significant reorganization of work in generations. Even if, in the long run, new jobs are created, I am aware that the timing of costs and benefits of AI may differ. Specifically, AI-related job loss could precede job gains. Although we do not have conclusive evidence of this occurring yet, it may still be on the horizon, and increased churn in the labor market could be anticipated.

Businesses are adopting AI at an increasing rate, but many have not yet used it to change the way they organize work. Indeed, the vast majority of small business respondents in the Federal Reserve's 2025 Small Business Credit Survey say that their labor costs have not changed as a result of AI.³ Yet, many businesses I hear from expect that AI will lead them to fundamentally change their business practices in the future.

Finally, I will turn to economic growth. Here, I am optimistic. Over the past year, gross domestic product (GDP) growth has remained robust. Labor productivity growth has exceeded its pre-pandemic average. I do not need to report this in the middle of Silicon Valley, but business creation has remained high. Having done research on innovation and its macroeconomic effects for the 20 years before I came to the Fed, I

³ See Federal Reserve Banks (2026), "2026 Main Street Metrics: Trends over Time from the Small Business Credit Survey," March 23, <https://doi.org/10.55350/sbcs-20260323>.

believe AI is a technology like none other I have seen in my lifetime. But as a student of Paul Romer and endogenous growth more generally, I have been waiting for this moment when the post-WWII investment in the knowledge economy would increase the arrival rate of ideas. As firms incorporate AI more systematically into their production processes, I expect that AI will further boost productivity growth, contributing to my expectation that GDP will grow robustly in the near to medium term.

What does this mean for monetary policy? I see elevated risks to both sides of our mandate, and from a risk-management perspective, I currently believe that the right course of action is to hold rates steady. However, I want to be clear about my risk assessment: The risks remain tilted toward higher inflation.⁴ In my baseline forecast, disinflation should resume in upcoming months without having to raise rates. Similarly, I expect the labor market will remain stable without having to lower rates.

After five years of above-target inflation, I am particularly attuned to the risk that elevated inflation will become embedded in price- and wage-setting behavior. As such, I am prepared to raise rates, if the expected disinflation does not appear in a timely manner. Likewise, I will continue to monitor labor-market developments, as well, and would be prepared to adjust my policy stance downward should the labor market deteriorate.

Responsibly Supporting Innovation

I will now turn to the theme of this conference: AI's effect on innovation, resilience, and risk in the financial system. I am excited to discuss this topic for at least two reasons. First, as an economist who began studying the economics of innovation in

⁴ See Lisa D. Cook (2026), "Economic Outlook," speech delivered at the Economic Club of Miami, Miami, Florida, February 4, <https://www.federalreserve.gov/newsevents/speech/cook20260204a.htm>.

earnest on this campus some time ago, I see great benefit that could come to the financial system from AI. Second, I serve as the chair of the Federal Reserve Board's Committee on Financial Stability. As a policymaker, figuring out how to encourage innovation, while ensuring the risks are contained and the system remains resilient, is a major preoccupation.

Overall, I would like to stress that I believe in experimentation. This approach thrives in Silicon Valley, and we are embracing it at the Fed, as well. That is why I cofounded the EmergingTech Economic Research Network, the System-wide effort to share AI research and results of AI experiments, and also why I have been encouraged to see staff at the Fed looking for ways to adopt AI technology in new and imaginative ways over the past several years.⁵ Not every effort will be a success, and, as I learned from my training here and at Berkeley, that is okay. We are seeing results from this experimentation-driven mindset, which I will talk about shortly after I discuss the broader benefits of AI-driven innovation in the financial system.

Benefits to Financial System of AI

I am optimistic about AI's promise to boost productivity and increase the arrival rate of ideas, which will support growth, lead to the creation of new firms that will produce new jobs, and put downward pressure on inflation. Within the financial sector, specifically, I am excited about the benefits from AI we are starting to see.

The financial sector is adapting to the current generation of AI tools and increasing its adoption, initially in highly manual or resource-intensive areas. This

⁵ Further information regarding the EmergingTech Economic Research Network is available on the Federal Reserve Bank of San Francisco's website at <https://www.frbsf.org/research-and-insights/emerging-tech-economic-research-network>.

transition includes in compliance functions, call centers, and back-office operations. Generating novel analytics has also become faster and more flexible. Using AI as a coding tool is helping the financial sector tackle age-old problems, such as updating legacy code and integrating systems. Next-generation models should more broadly adopt and integrate into client- and market-facing applications. Large technology and financial services firms, those who provide the hardware, software, and systems that underlie much of the global economy, use advanced AI tools to scan for potential cyber vulnerabilities that could be exploited. Further, AI adoption offers many opportunities for our financial system to be improved. These tools could allow firms to improve access to credit, allocate capital more efficiently, and speed processes. For example, AI could enable firms to accomplish the following:

- Develop new and better products that are more customized to individuals, broadening access to sophisticated financial products.
- Provide retail investors with the tools necessary to identify trends and emerging risks earlier.
- And leverage the benefits of efficiency gains to allocate more capital to lending and investments, which could lead to more economic activity and growth, as I mentioned earlier.

Risks and Vulnerabilities Related to AI

Broadly, I see AI as stimulating economic growth, which all else equal, should support financial stability. However, as a policymaker, I understand that innovation can lead to increased risk, if not monitored appropriately. I think about this likelihood both

through the lens of AI's interactions with long-standing vulnerabilities and of the risk a hypothetical AI shock would present to the system.

AI might introduce vulnerabilities to the financial system through a number of channels. One of the most commonly cited is the increased prevalence of AI-driven algorithmic trading. Traditional algorithms are fast, simple rules operating at nanosecond frequencies, but they are relatively rigid and hard coded. Generative AI and machine learning add self-learning based on historical experience, adaptation based on current market conditions, and analysis of unstructured data, such as text. Policymakers and academics have noted that, increasingly, AI-driven algorithmic trading may generate financial-stability risks, such as more correlated trading, endogenous model collusion, potential market manipulation, and greater market concentration.

Another potential risk comes from the probability that AI may displace or disrupt entire sectors. For example, concerns about AI disruption risk have affected speculative-grade bonds in the technology sector, where spreads have increased, as our *Financial Stability Report* noted earlier this month.⁶ These trends reflect AI disruption concerns in the software industry and arose after a large AI firm introduced products aimed at that sector. Concerns about credit exposure to software also contributed to the wave of redemptions that have put significant pressure on both traded and nontraded perpetual business development companies in recent months.⁷

⁶ See Board of Governors of the Federal Reserve System (2026), *Financial Stability Report* (Washington: Board of Governors, May), <https://www.federalreserve.gov/publications/files/financial-stability-report-20260508.pdf>.

⁷ See Paula Seligson, Olivia Fishlow, Rene Ismail, Davide Scigliuzzo, and Laura Benitez (2026), "Private Credit's Gate-Crashers Are Forcing Funds into a Brutal Spot," *Bloomberg*, March 8.

Another emerging trend that may have implications for financial stability relates to the fact that firms are increasingly tapping debt markets to finance the capital investments relating to AI infrastructure. Many of the hyperscaler firms have executed large investment-grade bond deals in recent months to fund AI capital expenditures.⁸ In addition, smaller data-center developers are raising debt from private debt funds, as well as asset-backed credit markets, to fund their investments.⁹ While many of the largest investors are also strong borrowers, the increasing use of leverage to finance investments in an emerging technology carries risk, and a sustained boom in debt issuance could eventually represent a financial-stability concern. I will note that even under very ambitious investment and debt-issuance projections, we would be unlikely to return to peak leverage levels observed before the Global Financial Crisis.

Cyber Risk

And, of course, we cannot talk about risk without discussing cyber risk. Recent advances in the ability of large language models (LLMs) and agentic AI systems to detect, exploit, and create new vulnerabilities have introduced new challenges in safeguarding system security for financial institutions, infrastructure, and third-party service providers. Very powerful AI tools, such as Anthropic’s Mythos Preview model, have demonstrated the ability to detect previously undetectable vulnerabilities in software applications that support important and widely used computer systems.

⁸ See Anhata Rooprai, Zaheer Kachwala, and Johann M. Cherian (2025), “Tech Companies Tap Debt Markets to Fund AI and Cloud Expansion,” *Reuters*, November 24 (updated May 11, 2026), <https://www.reuters.com/business/media-telecom/tech-companies-tap-debt-markets-fund-ai-cloud-expansion-2026-05-11>.

⁹ See Paula Seligson (2026), “The \$3 Trillion AI Data Center Build-Out Spurs a Debt Market Boom,” *Bloomberg*, February 2.

Non-malicious cyber events, such as software malfunctions, have also caused disruptions to the provision of financial services. AI can make developing software—particularly writing code—faster and easier. However, by contributing to the rapid proliferation of code, the aggressive use of AI may indirectly strain current security review processes.

The ultimate implications of AI for cybersecurity remain unclear. Advanced AI coding agents can be used to enhance the security of many important computer systems to prevent future AI-related cyberattacks.¹⁰ It remains possible that AI makes financial institutions more resilient regarding cyberattack vulnerabilities.

AI at the Fed

Just like financial firms and other entities across the economy, the Fed is also working to responsibly deploy AI to advance our mission and to improve our own work. To be clear, as I said at the NBER Summer Institute last year, the Federal Open Market Committee is not using AI in developing or setting policy¹¹. But many parts of the Fed system are using AI for a variety of other tasks, particularly in the area of financial stability, and we already see tangible benefits. By using AI ourselves, we can improve our analysis of the financial sector and are better able to highlight vulnerabilities—whether they are new ones introduced by AI or old ones that we may have missed. The use of AI can make us better at our job with enhanced monitoring and improved analysis.

¹⁰ See Saeed Azhar, Tatiana Bautzer, Michelle Price, and Francesco Canepa (2026), “Anthropic’s Mythos Sends U.S. Banks Rushing to Plug Cyber Holes,” *Reuters*, May 12, <https://www.reuters.com/business/finance/anthropics-mythos-sends-us-banks-rushing-plug-cyber-holes-2026-05-12>.

¹¹ See Lisa D. Cook (2025), “AI: A Fed Policymaker’s View,” speech delivered at the National Bureau of Economic Research, Summer Institute 2025: Digital Economics and Artificial Intelligence, Cambridge, Mass., July 17, <https://www.federalreserve.gov/newsevents/speech/cook20250717a.htm>.

I would like to share two specific ways that we are using AI to advance our critical mission of monitoring financial stability.

First, newly formed teams of experts within the Division of Financial Stability are analyzing technological risks to financial stability. These collaborative groups assess how cyber, AI, and quantum computing create both vulnerabilities and opportunities. For example, economists Anne Lundgaard Hansen at the Richmond Fed and Seung Jung Lee at the Board have investigated the effect of generative AI adoption on financial stability through laboratory-style experiments using LLMs.¹² Their research on herd behavior in investment decisions in a stylized lab setting found that AI agents make more rational decisions than humans. The research suggests that agents are more likely to make decisions based on data and analysis, rather than simply following general market trends. This outcome could potentially lead to fewer asset price bubbles arising from animal spirits.

These innovative teams have also developed practical tools for our mission. One team designed a method to construct a small, cost-efficient AI model that can classify a large amount of text just as accurately as a larger model, using a technique called “active knowledge distillation.” The method achieves up to an 80 percent reduction in computation costs while maintaining accuracy.¹³ This efficiency matters, because financial-stability analysis increasingly requires processing vast amounts of unstructured text data, including regulatory filings, earnings calls, and news articles. Another

¹² See Anne Lundgaard Hansen and Seung Jung Lee (2025), “Financial Stability Implications of Generative AI: Taming the Animal Spirits,” Finance and Economics Discussion Series 2025-090 (Washington: Board of Governors of the Federal Reserve System, September), <https://doi.org/10.17016/FEDS.2025.090>.

¹³ See Viviana Luccioli, Rithika Iyengar, Ryan Panley, Flora Haberkorn, Xiaoyu Ge, Leland Crane, Nitish Sinha, and Seung Jung Lee (2025), “LLM on a Budget: Active Knowledge Distillation for Efficient Classification of Large Text Corpora,” Finance and Economics Discussion Series 2025-108 (Washington: Board of Governors of the Federal Reserve System, December), <https://doi.org/10.17016/FEDS.2025.108>.

interesting project applied natural language processing to decades of Beige Book data, finding that even when controlling for traditional metrics, the sentiment in these anecdotal compilations provides meaningful explanatory power in forecasting recessions.¹⁴

Second, the staff from the Board and all 12 Reserve Banks recently participated in an agentic AI sprint. This event encouraged experimentation and explored what agentic AI could do for financial-stability analysis. It was great to see all the AI systems that could reason through problems, decide which analytical approaches to use, and complete complex tasks with minimal human intervention. A valuable insight we found in one of the projects was that agentic AI systems can be more systematic in identifying network-based risks than our standard approaches. This outcome was not because we do not understand their theoretical importance, but because, in many cases, we lack the capacity to comprehensively analyze complex empirical structural patterns of networks at scale.

This type of systematic capability translates into potentially meaningful efficiency gains for financial-stability work. For example, other prototypes demonstrated that they could select, run, and analyze many financial-stability–relevant scenarios that would be prohibitively time-consuming using traditional methods. This development enables the kind of thorough analysis that humans would struggle to complete in a reasonable time frame. However—and this is critical—systematic coverage without accuracy would be worse than a selective approach. The most promising approaches build verification into the system architecture itself. These approaches have multiple agents confer before

¹⁴ See Shengwu Du, Flora Haberkorn, Isabel Kitschelt, Seung Jung Lee, Anderson Monken, Dylan Saez, Kelsey Shipman, and Sandeep Thakur (2026), “Do Anecdotes Matter? Exploring the Beige Book through Textual Analysis from 1970 to 2025,” Finance and Economics Discussion Series 2026-004 (Washington: Board of Governors of the Federal Reserve System, January), <https://doi.org/10.17016/FEDS.2026.004>.

reaching a consensus and include mechanisms that force the agents to consider contrarian perspectives. This process, in turn, can be crosschecked by researchers. If that sounds familiar at a place like Stanford, it should. What I just described for AI agents are the same types of methods that have yielded breakthrough thinking from humans for centuries.

Conclusion

The totality of our experience with AI leads us to the conclusion that, alongside experimentation, strong governance and risk management must be our foundation. The most promising approaches augment human judgment with AI capabilities while building verification into the architecture itself.

The urgency is real. AI is advancing rapidly, and financial institutions are adopting these technologies apace. As policymakers, we must understand these systems through hands-on experience. By building our own AI capabilities, we gain invaluable insights into both the promise and risks these technologies bring to the financial system. With appropriate governance frameworks, autonomous intelligence can potentially expand our analytical capabilities. These tools could enhance our capacity to identify and respond to evolving threats. But we proceed with both optimism and caution, as warranted at this moment of a technological inflection point.

Thank you again for the opportunity to return to Stanford and to speak at this timely and important event. I look forward to our discussion.