

Artificial intelligence and Financial Stability

As prepared for delivery, ECB Sintra Forum on Central Banking, June 30, 2026

Tobias Adrian

Thank you very much. It is a pleasure to join this panel on artificial intelligence and financial stability, and I am grateful for the opportunity to contribute to this timely discussion.

I would like to organize my remarks around two broad questions. The first is whether the rapid rise in valuations of artificial intelligence firms should be viewed as a source of systemic risk. The second is how artificial intelligence is changing the practice of financial intermediation—and what that implies for policymakers, regulators, and supervisors.

AI valuations and financial stability

Let me begin with valuations. Much of the recent discussion has focused on the strong performance of technology firms, particularly in the United States, where many of the leading AI companies and infrastructure providers are located. Looking at broad equity markets, such as the S&P 500, valuation pressures had reached very elevated levels earlier this year—around the 95th percentile by historical standards, and on some days even higher. More recently, those pressures have moderated, moving closer to the 80th percentile.

That moderation is important. It reflects some adjustment in prices, but also the fact that earnings have continued to surprise on the upside. In other words, the rise in equity prices has not been driven by valuation dynamics alone--it has been supported by very strong expectations and even stronger realizations of profitability, particularly in the parts of the technology sector linked to AI chip and frontier models.

This is one reason why the current episode does not look, at least so far, like a classic asset-price bubble. In a bubble, one would expect a broad weakening of market discipline, with prices rising indiscriminately across the sector. By contrast, market forces still appear to be differentiating across firms. Some chipmakers have continued to see very strong valuations, reflecting demand for AI infrastructure. At the same time, parts of the software sector have sold off sharply, in part because AI itself may disrupt existing business models.

That said, there are financial stability issues that deserve close attention. One concern is that some of the largest technology firms and hyperscalers are beginning to increase leverage as they invest heavily in AI infrastructure and finance about half of that in credit markets. These investments include physical assets—such as chips and data centers—that may depreciate or become obsolete relatively quickly, while the financing may be medium- or long-term. This creates a potential mismatch between the economic life of the physical assets and the maturity of the liabilities.

As long as profitability remains strong, and as long as firms and households continue to pay for frontier AI services at scale, that mismatch may be manageable. But if earnings disappoint—whether for the sector as a whole or for individual firms—financial stability risks could become more salient. The key issue, therefore, is not simply whether valuations are high. It is whether the profitability underpinning those valuations proves durable.

And there are a number of reasons to worry about profitability. One is the ability to charge for models, given competition from open source models. Another is the competition among frontier large language model providers, which may erode profitability for some of those providers at some point.

How AI is changing financial intermediation

Let me turn to the second question: how AI is changing the financial sector itself. I will focus on three areas—lending, trading, and risk management.

Lending

Financial institutions have used quantitative methods in lending for many years. Credit scoring, statistical models including machine learning tools, and large data sets are already embedded in many parts of consumer and corporate credit markets.

Generative AI, however, may represent a more transformative shift. When we speak with banks and other market participants, they increasingly describe AI as a technology that augments the capacity of loan officers and credit teams. In the same way that AI can help an economist review a large literature or summarize a long record of speeches, generative AI can help a loan officer process information, identify patterns, and form judgments more efficiently.

This augmentation could raise productivity and improve access to credit. It may also enable models to identify lending opportunities that were previously difficult to detect. Over time, more lending decisions could become increasingly automated, or agentic, with AI systems playing a more direct role in the origination and monitoring of credit. That is even more so to the extent that lending decisions may become

“agentic”, meaning that generative AI models would evolve the nature of the decisions process autonomously.

This also creates supervisory challenges. The central question is how these systems perform in a credit downturn. Models that appear effective in normal times may behave differently when defaults rise, liquidity tightens, or borrower behavior changes. In addition, many generative AI systems remain difficult to explain. If a loan is approved or denied, supervisors—and firms themselves—need to understand the basis for that decision, especially when models are complex, dynamic, or partially opaque.

Standard setters and regulators have already begun to issue guidance in this area. But the real test will come through the cycle. We will need to see how AI-assisted and AI-driven lending decisions behave under stress, and whether governance, validation, and accountability frameworks are sufficiently robust.

Trading

The second area is trading. Automated trading is not new. In many liquid markets—equities, foreign exchange, and on-the-run sovereign bonds—large shares of trading volume are already driven by models, including high-frequency trading strategies that are heavily using machine learning.

Generative AI may accelerate this process by changing how trading strategies are developed, tested, and deployed. Coding can be faster. Back-testing can be faster. Strategies can move from design to market more quickly, and can also be withdrawn more quickly when they do not perform as expected.

There are clear efficiency gains from this. AI can also help institutions upgrade enterprise platforms, improve data infrastructure, and strengthen operational processes. But there is also a potential financial stability concern: the speed and automation that already characterize liquid markets may migrate to less liquid asset classes.

In a downside scenario, this could matter. Less liquid, less centralized, and over-the-counter markets may not have the same safeguards as exchange-traded markets. Circuit breakers, kill switches, and other tools that have helped contain stress in highly liquid markets may be less effective in private markets or more fragmented trading environments. The policy question is therefore whether our existing tools are well suited for a world in which AI-enabled trading becomes faster, broader, and more agentic.

When trading systems are entirely agentic, and interact among each other in the marketplace, novel market conduct questions arise. Agentic trading may promote collusion and increase systemic risk via strategic interactions. Such developments may be difficult for central banks and regulators to detect, as trading strategies may evolve rapidly, be opaque, and interactions may be difficult to monitor in real time.

Risk management

The third area is risk management. Here, AI offers important opportunities. It can help firms detect anomalous employee behavior, identify fraud, monitor cyber vulnerabilities, and analyze complex trading books. It may help risk managers see patterns that would otherwise be missed, especially in large and complex data sets.

At the same time, AI also expands the toolkit available to bad actors. Fraud attempts may become more sophisticated. Cyberattacks may become more targeted. Social engineering may become more convincing. The most sophisticated firms may be relatively well positioned to respond because they have access to frontier models, advanced infrastructure, and specialized human capital.

But financial stability is not only about the most sophisticated firms. It is also about the resilience of the broader system, including smaller institutions and firms in jurisdictions with more limited technological capacity. Attackers often look for the weakest link. If AI increases the sophistication of threats, then uneven defensive capabilities across the financial system could become a source of vulnerability.

Policy implications

Let me conclude with three policy implications.

First, policymakers need to understand what augmentation means for decision-making. AI may improve productivity and expand capacity, but it also changes accountability. When human judgment is supported—or partially replaced—by AI systems, firms and supervisors need clarity about who is responsible for outcomes.

Second, we need to pay close attention to speed. AI can compress the time between information, decision, and action. That may improve market efficiency in normal times, but it may also amplify stress in downside scenarios, particularly in less liquid markets where traditional safeguards are weaker.

Third, we need to address the explainability gap. Agentic lending and trading can make risk taking difficult to understand, particularly at the system wide level. Supervisors do not need every model to be perfectly transparent in every technical detail. But they do need enough understanding to assess governance, risk controls,

fairness, robustness, and resilience. Without that, accountability becomes difficult and trust in financial intermediation may weaken. But implementation in the face of agentic decision making may challenge those principals.

In sum

Artificial intelligence offers substantial opportunities for the financial sector. It can raise productivity, improve risk management, and support better decision-making. But it also creates new channels through which vulnerabilities can build and propagate. The task for policymakers is not to slow innovation, but to ensure that innovation is accompanied by strong governance, effective supervision, and a financial system that remains resilient through the cycle.

Thank you.