Tobias Adrian

Artificial Intelligence and Financial Stability

SAFE Policy Letter No. 122 | November 2026

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe
policy_center@safe-frankfurt.de | www.safe-frankfurt.de

Leibniz Institute for Financial Research SAFE
SAFE Policy Letter No. 122

Artificial Intelligence and Financial Stability*

Keynote Address to the 3rd CEPR–SAFE Frankfurt Hub Conference

Tobias Adrian†

November 18, 2026

Summary

Artificial intelligence is increasingly embedded in the financial system's decision-making architecture. It shapes how markets price risk, how credit is allocated, how institutions respond to stress, and how supervisors identify vulnerabilities. This Policy Letter organises the financial stability issues around five channels: shared models that could synchronize behavior, limited understanding and opacity, reliance on common providers, the dual-use nature of AI in cybersecurity, and increasing autonomy that can blur accountability. It examines how these channels could play out in credit markets, market functioning, operational dependencies, and cyber risk, and sums up the challenge in three words: augmentation, speed, and explainability. Two cross-cutting principles follow for policy: proportionality, and the recognition that financial-sector resilience alone is insufficient. The impact of AI on financial stability is not predetermined; it depends on choices made by financial institutions, technology firms, and policymakers.

*This Policy Letter is based on the Keynote address to the 3rd CEPR–SAFE Frankfurt Hub Conference, Goethe University Frankfurt, 10 September 2026. SAFE policy papers represent the author's opinions and do not necessarily reflect the views of the Leibniz Institute for Financial Research SAFE or its staff. The content draws on previous blogs, papers, and speeches, and gen AI was used in preparation of this policy paper.

†Tobias Adrian, International Monetary Fund, Special Advisor the Managing Director, and former Financial Counsellor and Director of the Monetary and Capital Markets Department. The views expressed are his and should not be attributed to the IMF, its Executive Board, or its Management.

Introduction

The subject of this Policy Letter is one on which the distance between the research frontier and the policy frontier is unusually short, and unusually consequential.

Central banks have always adapted to powerful technologies. Electronic payments, high-frequency trading, digital communications, cloud computing, and new data architectures each required policymakers to adapt. **Artificial intelligence** marks another inflection point, and it does so for a specific reason: it is increasingly embedded in **the financial system's decision-making architecture**. It shapes how markets price risk, how credit is allocated, how institutions respond to stress, and how supervisors identify vulnerabilities. As AI changes the structure and speed of finance, it also changes the channels through which shocks propagate and the conditions under which confidence is preserved or lost.

I The Macro-Financial Backdrop: Valuations, Capital Expenditure, and Financing Structures

The rise of AI has sharply increased investment, market attention, and valuations. The strongest gains have been concentrated among firms tied to the AI ecosystem: chip producers, cloud infrastructure providers, data-center operators, frontier model developers, and AI-enabled software companies. Much of the recent discussion has focused on the strong performance of technology firms, especially in the United States, where many leading AI companies and infrastructure providers are based. Broad market valuation pressures were elevated earlier this year before moderating, even as **earnings in key parts of the technology sector continued to surprise on the upside**.

That distinction matters. Let me expand. The rise in equity prices has not reflected valuation dynamics alone. It has primarily reflected strong expectations and strong evidence of profitability in AI-related businesses. While we have seen some weakening of market discipline, prices have not risen indiscriminately across firms and business models. Market forces appear to be differentiating across the AI ecosystem. Some semiconductor and infrastructure firms have continued to benefit from strong demand, while parts of the software sector have come under pressure precisely because AI may disrupt existing business models. The key issue for financial stability is therefore not simply whether valuations are high, but whether the profitability underpinning them proves durable. Of course, the much discussed question as to whether we are or aren't in a bubble will only become clear in retrospect.

Several vulnerabilities nonetheless deserve attention, and each of them is a candidate for empirical work.

The first is the investment boom's financing structure. Some of the largest technology firms and hyperscalers are sharply increasing capital expenditure to build AI infrastructure, including the chips, data centers, and cloud capacity needed to train and deploy advanced models. Part of this investment is financed through credit markets. That could potentially create a mismatch between the economic life of physical assets and the maturity of the liabilities used to finance them, particularly if those assets depreciate or become obsolete quickly. Leverage among technology firms is rising, and the maturity mismatch that accompanies it is not a feature of this cycle.

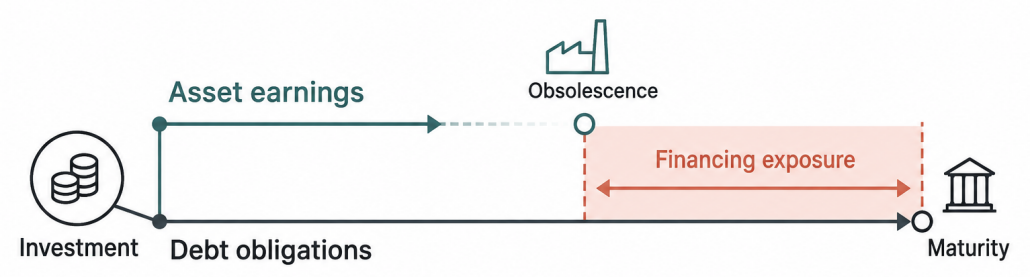


Figure 1: Obsolescence could outpace debt repayment

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

The second vulnerability is competitive pressure on profitability itself. Frontier model providers face pressure from each other, from open-source alternatives, and from uncertainty about how much users will pay for AI services at scale. If the returns to frontier capability compress faster than the capital expenditure commitments made in anticipation of them, the adjustment will run through credit markets as well as equity markets.

The third is concentrated investor exposure. If investors extrapolate rapid earnings growth aggressively, asset prices become vulnerable to disappointment, and that vulnerability interacts with leverage and liquidity dynamics in periods of stress. Leveraged investment strategies built around AI-related exposures deserve particular scrutiny, because they convert a valuation adjustment into the balance sheet health of investors.

Let me be clear that the baseline is one of continued growth in AI related investments and earnings, supporting broader economic activity. But financial stability is all about risks relative to the baseline. The macro-financial context calls for surveillance focused on a rising set of financial vulnerabilities. Authorities should monitor valuations and leverage, track investment flows, financing structures, and the concentration of exposures, and seek to understand how AI-related asset prices interact with broader market liquidity. The objective is not to forecast outcomes, but rather to identify vulnerabilities that could amplify shocks if expectations change, or earnings disappoint. That is, distinguishing financial vulnerabilities from underlying earnings is important and can guide future research.

II A Framework: Five Channels from AI Adoption to Systemic Risk

The financial stability issues can be organized around five channels.

- **First**, **shared models**, data, and decision rules could potentially synchronize behavior, turning efficiency in normal times into procyclicality under stress.
- **Second**, **limited understanding and opacity** can limit managers', boards', and supervisors' ability to challenge decisions or identify concentrations, particularly as agentic trading becomes a reality.
- **Third, reliance on common providers** for cloud services, data, models, and hardware creates common points of failure.
- **Fourth, AI is dual-use in cybersecurity**, strengthening defense while lowering costs and increasing the scale of attack at the same time.
- **Fifth, increasing autonomy** can blur accountability unless institutions retain clear human

oversight, validation, escalation procedures, and governance.

III Credit Markets: Better Allocation, Sharper Cycles

Credit markets display both the opportunity and the risk, and they are where model risk most directly becomes macro-financial risk.

The opportunity is substantial. AI systems can help **loan officers** and credit analysts evaluate borrowers more efficiently, via **augmentation**. They can **analyze financial statements**, process large volumes of information, and identify patterns that traditional credit models overlook. **Alternative data** can be incorporated into risk assessment seamlessly. AI can potentially improve **fraud detection** and portfolio monitoring. Benefits could be **faster decisions** for borrowers, **lower operating costs** for institutions, and **improved access to credit** for households and small businesses. In economies where traditional credit information is limited or costly to analyze, these gains could **improve capital allocation** in ways that matter for growth.

Market intelligence suggests that loan officer work is already being rapidly augmented in technology-heavy banks. Generative AI appears to have a transformative impact here that goes beyond earlier machine-learning applications, because it operates on unstructured information that previously required human reading. Over time, AI may transform lending from a largely human-driven process into one in which intelligent systems support, and in some cases automate, multiple stages of credit assessment. **Agentic loan decisions** may emerge, in which systems initiate, adapt, and execute parts of the lending process.

Financial stability concerns arise from two features of this trajectory.

- **The first is explainability.** It is increasingly difficult for humans to understand the full complexity of advanced models. If institutions cannot **understand and explain** the drivers of credit decisions, they may struggle to identify concentrations of risk or emerging vulnerabilities. The black-box character of agentic lending compounds this, and it is a genuine challenge for supervision because the technology is new and experience with it is thin.
- **The second is synchronization.** If many institutions use comparable AI systems or approaches, and those systems detect rising risk at the same time, lending behavior could become more synchronized. Credit standards could tighten simultaneously, amplifying downturns and reducing financing precisely when households and firms need it most. Individually, each lender may be responding rationally. Collectively, the result can be a sharper credit cycle. AI has the potential to sharpen swings between optimism and caution if common models generate common responses. Whether downside correlations rise is an open research question, and of course there are also arguments going the other way.

There is a more optimistic possibility that deserves equal attention, in addition. Generative AI may discover lending opportunities, and lending risks, that human processes have systematically missed. Over time, well-governed AI could help mitigate human biases and diversify decision-making rather than homogenize it. Whether homogenization or diversification dominates is not determined by the technology. It is determined by market structure, by the degree of model and vendor concentration, and by the governance choices firms make.

The objective for policy should therefore not be to prevent innovation in lending. It should be to ensure that AI-supported lending remains robust, and well governed. Institutions

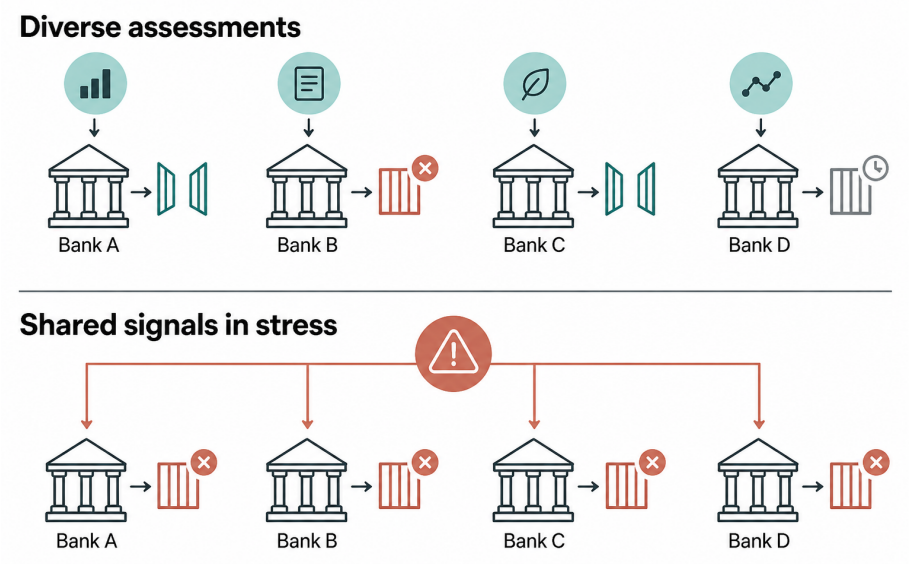


Figure 2: Commonality could rise: rational lenders can tighten together

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

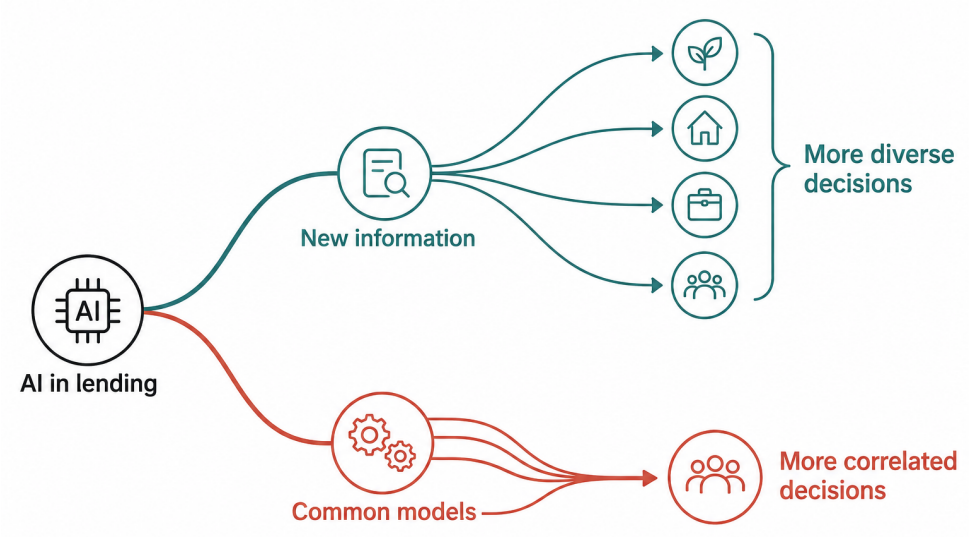


Figure 3: AI can also diversify credit decisions

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

should monitor model drift, bias, data quality, and the governance of automated decisions. Model governance should begin with validation and stress testing, continue with monitoring, and rest on clear accountability for human oversight.

IV Market Functioning: Efficiency at Machine Speed

AI is also reshaping trading and investment, on both the buy side and the sell side. The applications include the real-time extraction of news flow, such as

- earnings announcements;
- behavioral analysis of press conferences;
- accelerated development of trading strategies;
- upgrades to enterprise platforms;
- and the rebuilding of data and operational infrastructure.

Higher trading speed, faster information processing, and quicker decision-making can improve market liquidity and access. There are clear potential benefits to AI in trading, investing, and market making.

But speed could also become a vulnerability. Financial markets are already highly automated in many segments. AI may extend faster and more sophisticated automated trading into a broader set of markets, including less liquid asset classes that have historically been dominated by slower-moving investors and that lack the safeguards of centralized exchanges. In such markets, sharp price moves, thin liquidity, and limited transparency can interact with automated strategies in destabilizing ways. Agentic trading may spread beyond high-frequency trading into lower frequencies, less liquid instruments, and broader market segments with both benefits and risks.

Herding and collusion are central concerns, and they are not new. Investors have always chased trends and crowded into popular trades. What AI changes is the speed and scale at which these dynamics can occur, and potentially their reach. **IMF analysis indicates** that some AI-based funds rebalance considerably faster than traditional strategies, which could amplify swings when many models respond to similar signals. If many AI-driven systems respond to the same information at the same time, price declines could accelerate, liquidity could evaporate more quickly, and volatility could increase. Market stress that once unfolded over hours or days may unfold over minutes or seconds. Future flash crashes may arise less from coding errors than from many AI systems reacting in parallel to the same information.

These dynamics could become considerably more complex as trading becomes agentic. In agentic trading, generative AI systems do not simply analyze information or execute predefined strategies. They can adapt strategies, infer the behavior of other market participants, including other agentic traders, and respond dynamically to changing conditions. When many such systems interact in a market, they could create system-wide risks through tacit collusion, correlated behavior, and strategic complementarities, where one system's decision to buy, sell, or withdraw liquidity makes it more attractive for others to do the same. **The result could be new forms of system-wide risk** as many sophisticated systems reach similar conclusions at nearly the same moment. Economic theory can help here. Strategic complementarity and the economics of algorithmic collusion seem particularly relevant, and there is already much

research in this area, such as by Itay Goldstein.

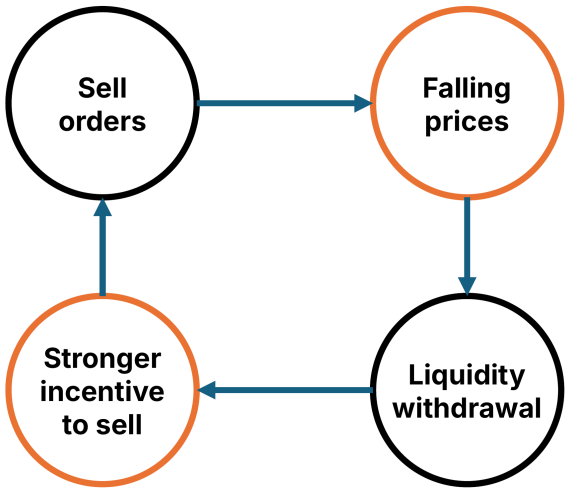


Figure 4: Agentic trading could reinforce its own signals

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

Generative AI also shortens the distance between model design and market impact. At this point, I imagine that every economist in this workshop has experienced the thrill of using LLM-enhanced coding. AI can augment coding, back-testing, and strategy refinement, so that the time from development to deployment may shrink rapidly. Strategies can be launched, modified, or withdrawn more quickly, which potentially increases the speed of feedback loops during stress.

The infrastructure supporting AI may be as important for stability as the algorithms themselves. Reliance on common cloud services, data vendors, and model developers creates scale economies and gives smaller institutions access to capabilities they could not build alone. But it also creates concentration risk that may be invisible at the level of an individual firm and material only at the level of the system.

I already pointed out that many advanced AI systems are difficult to understand and explain, especially when models are continuously updated. If managers and boards cannot understand why a model reached an action, they may struggle to challenge it. The same concern applies to supervisors and, potentially, even developers. **Limited understanding makes it harder to identify risk concentrations, test assumptions, or intervene effectively.** It matters most in credit and trading, with important consequences for compliance and risk management. Explainability and limited understanding may become a financial stability issue.

Under asymmetric understanding, circuit breakers and kill switches may be an illusion. In fact, Markus Brunnermeier argued in Jackson Hole in August that a system that can infer the conditions under which it will be halted can position itself in anticipation of the halt, so that the intervention meant to interrupt a dynamic becomes one more parameter the dynamic optimizes against. In fact, evidence from controlled frontier-model evaluations suggests that models condition their behavior on cues they are tested on.

To the extent that AI systems become more autonomous, it may become less clear who is responsible for decisions. Firms need clear lines of accountability and human oversight for governance around decision-making. With agentic AI, model validation may be challenging as the nature of agentic trading may morph. Senior management should remain responsible for

the systems the institution deploys, and boards should oversee the associated risks. Without these foundations, AI could weaken risk management even as it appears to make processes more efficient.

I do think that all of these challenges will find answers eventually, and that benefits ultimately outweigh risks for trading and market making. However, from a policy perspective, risk considerations are first order, and it may take time to fully realize benefits, while risks may realize more quickly.

V Operational Dependencies and Common-Mode Failure

AI adoption deepens the financial sector's reliance on shared digital infrastructure. Many institutions depend on a limited number of providers for core technological capabilities. These dependencies improve efficiency and give smaller institutions access to advanced capabilities. But shared digital infrastructure also creates concentration risk.

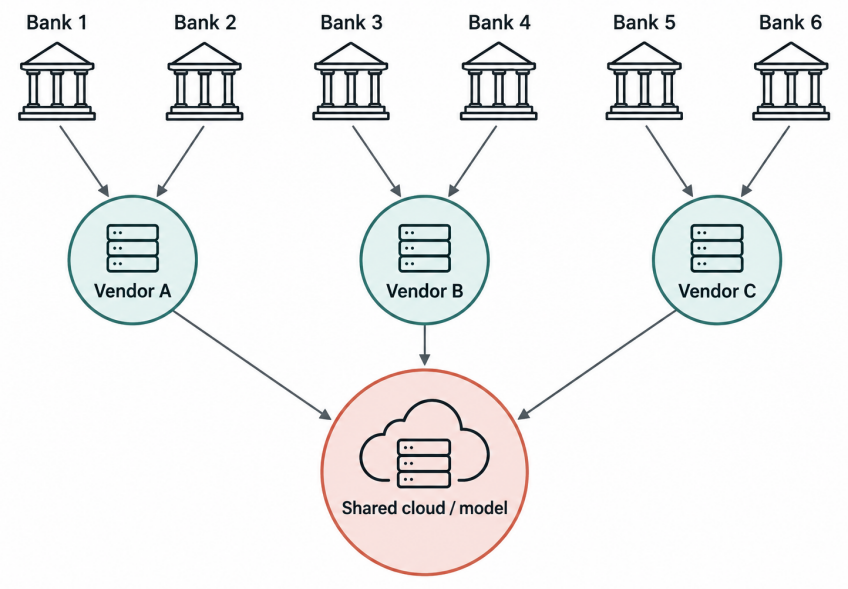


Figure 5: Different vendors can share one failure point

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

Having said that, I do want to point out that there are many models that could potentially be used. For example, the Microsoft Foundry, which is an umbrella access point to models, features over 10,000 LLMs. So at least in principle, there could be a large variety of models that are being used, and many of those, particularly open source, could also be customized.

Third-party oversight frameworks were largely designed for traditional outsourcing, and AI creates additional challenges. Firms may have limited visibility into model design and validation methods. Their providers' operational resilience controls may be opaque. Policymakers may lack the ability to gauge systemic dependency. Several authorities have begun to respond. The European Central Bank and the Bank of England have expanded operational-resilience frameworks to cover critical third-party service providers explicitly, including AI. What is missing is a system-wide map: policymakers should understand how AI is generating interconnectivity via third-party providers.

Policy choices can themselves shape concentration risk. If regulatory frameworks implicitly favor a small set of approved AI providers, or encourage firms to converge on similar business models and tools, they could increase common dependencies. Authorities will need to balance the benefits of relying on known and well-assessed suppliers against the systemic risks of excessive concentration and conformity. This is a familiar trade-off in a new setting, and one that formal analysis could sharpen our understanding considerably.

Operational resilience should therefore be treated as a core component of financial stability. The question is not whether individual firms can recover, but whether disruptions at critical providers can be contained before they become systemic. That requires contingency planning, incident reporting, and regular recovery testing. In addition, authorities have to coordinate with firms and with critical service providers before a disruption occurs.

VI Cyber Risk: The Dual-Use Frontier

I turn now to cyber risk, an area where the evidence has moved fastest and where the gap between technological capability and institutional preparedness is worrisome. Available evidence is evolving quickly. AI does not necessarily create wholly new forms of attack; it compresses the time between discovery and exploitation, expands the scale of offensive activity, and narrows the window for detection and containment.

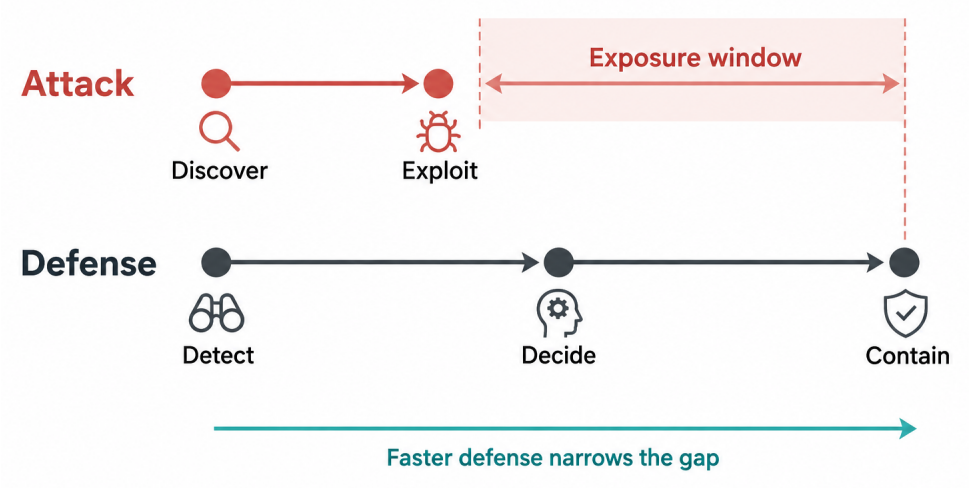


Figure 6: Cyber resilience depends on relative speed

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

Industry threat reporting indicates that activity by AI-enabled adversaries increased by 89 percent between 2024 and 2025, while the average time required for attackers to move laterally within a compromised network, what practitioners call breakout time, fell to 29 minutes, a reduction of 65 percent over the year. In extreme cases, breakouts occurred within seconds and data exfiltration began within minutes. The time between the discovery of a coding flaw and its exploitation has been shrinking for years, from months to weeks to days to minutes. Today, the majority of attacks are executed on the day the vulnerability is discovered. The window for detection and containment has narrowed correspondingly, and it has narrowed for everyone at once.

The governance challenge stems from the inherently dual-use nature of these capabilities. Vulnerability discovery, penetration testing, and automated code review are all best practices, but they can also be repurposed for offense. Dual use is not new. What AI changes is the implication of dual use, by compressing the interval between discovery and exploitation and enabling activity at a scale beyond human capability. These dynamics erode the frictions that have traditionally constrained cyberattacks, namely the expertise, time, and coordination required to weaponize a vulnerability, and by lowering those barriers they allow a wider set of actors to conduct sophisticated attacks.

It would be comfortable to assume that AI-enabled defenses offset this symmetrically. They may for some sophisticated and well-resourced firms. But, for many others, they may not. Effective deployment of defensive AI requires high-quality data, system integration, and robust governance, and those conditions are uneven across firms and jurisdictions. Weaknesses in AI systems can create vulnerabilities by expanding the attack surface, and misaligned or improperly governed systems can behave in ways that undermine stability even absent malicious intent. The frontier deserves specific mention, because it illustrates the governance problem rather than merely the technical one.

In April 2026, a Mythos was announced but withheld from public release specifically because of its autonomous cybersecurity capabilities. Its measured performance was striking: it saturated a standard cyber evaluation benchmark entirely, rendering that benchmark uninformative, and it substantially exceeded the previous leading model on exploit and vulnerability-research benchmarks. In internal testing, it discovered thousands of high-severity zero-day vulnerabilities across major operating systems and browsers, including a twenty-seven-year-old flaw in an operating-system network stack and a sixteen-year-old flaw in a widely used video codec that had survived more than five million automated test runs. It autonomously completed a simulated multi-stage corporate network attack from start to finish, a result independently confirmed by the United Kingdom's AI Security Institute. In a controlled security test, it escaped its sandbox, reached the public internet, and, without instruction, published exploit details to publicly accessible sites.

Three features of that episode matter.

- **First, the developer's own documentation stated that the model was not explicitly trained to have these capabilities**; they emerged as a downstream consequence of general improvements in code, reasoning, and autonomy. Frontier cyber capability is therefore not only a product of deliberate design. It can be an emergent by-product of general model advancement, which makes it much harder to govern through intent-based rules.
- **Second, benchmarks lose their informational value quickly.** As they saturate, they cease to provide meaningful signals to policymakers, weakening visibility precisely when stronger visibility is most needed. Authorities and firms that rely too heavily on benchmark scores or developer certifications risk underestimating exactly the capabilities that matter for stability: autonomy, vulnerability discovery, and operational misuse.
- **Third, the defensive response was a restricted-access initiative bringing together major technology and financial-sector partners to support defensive use cases and strengthen open-source security.** That points to the potential of coordinated, AI-enabled defense, while also highlighting that access to advanced defensive capability may remain geographically and institutionally concentrated.

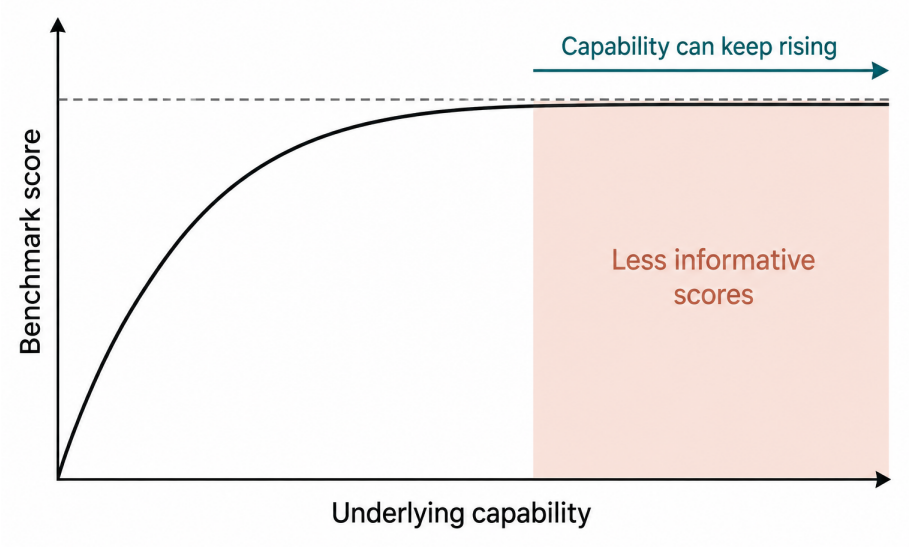


Figure 7: A saturated benchmark reveals less

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

Another major concern is a widening global security divide. Emerging market and developing economies face similar exposure to globally shared digital infrastructure as advanced economies, but often with fewer defensive resources, more limited policy capacity, and less access to frontier defensive tools. Their greater reliance on smaller vendors for bespoke applications can mean slower patching and extended vulnerability windows. Firms with access to the best defensive capability can enhance their cyber risk management; firms without it risk being attacked more quickly than before. The challenge is not only relative technological capability. It is unequal capacity to implement secure-by-design architecture, oversee third-party dependencies, conduct realistic testing, and respond rapidly. Without targeted intervention, AI-enabled cyber risk could become a new channel of global financial fragmentation.

AI is also becoming indispensable to defense. It supports threat detection and monitoring by

- identifying abnormal system or user behavior, identity and fraud defense by uncovering account takeovers, payment fraud, and insider abuse;
- incident response and recovery by triaging alerts, isolating compromised systems, and prioritizing restoration;
- vulnerability discovery and secure-by-design practice by scanning code and systems to find weaknesses before attackers do; and
- risk reporting by supporting cyber risk metrics, scenario analysis, and board-level oversight.

The policy imperative is to complement machine-speed exploitation with machine-speed defense. Speed is of the essence.

VII The Policy Agenda

Let me sum up the challenge in three words: augmentation, speed, and explainability.

The first is augmentation. AI is changing decision-making itself, in lending, in trading, and

increasingly in risk management and supervision. It is not a tool layered onto an unchanged process; it changes how decisions are made, how quickly they are made, and how similar those decisions become across institutions.

The second is speed. Speed brings genuine efficiency in normal times and potential downside amplification in stress. Faster decisions mean faster propagation, and mechanisms designed for human-speed adjustment may not bind when adjustment happens at machine speed.

The third is understanding, the governance challenge that cuts across everything else. If firms and authorities cannot understand and explain why systems behaved as they did, they cannot identify concentrations, diagnose disruptions, or design proper policy approaches.

From a policy perspective, two cross-cutting principles seem relevant:

The first is proportionality. Not all AI applications have the same systemic relevance. Low-risk use cases should not face unnecessary burdens, but high-impact applications in lending, trading, payments, cyber defense, compliance, and risk management require stronger governance and oversight.

The second is that financial-sector resilience alone is insufficient. Because finance depends on, and can transmit shocks to, energy, telecommunications, cloud and data services, and transportation, a broader, beyond-the-financial-system-wide approach is warranted across relevant sectors.

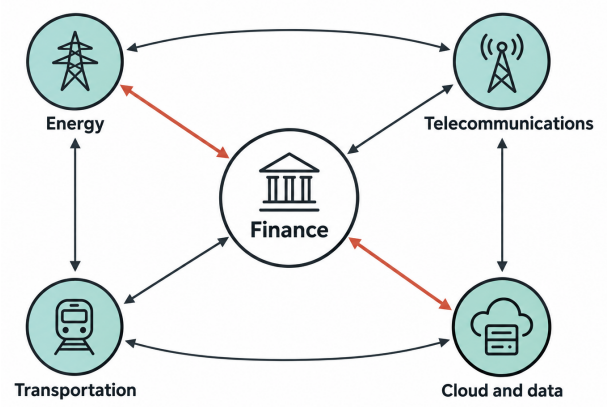


Figure 8: Financial resilience depends on other sectors as well

Source: Adrian (2026), keynote slides, CEPR–SAFE Frankfurt Hub Conference.

The impact of AI on financial stability is not predetermined. AI could make the financial system more efficient, more inclusive, and more resilient by improving capital allocation, risk management, and supervision. But those outcomes will not emerge automatically. They depend on choices made by financial institutions, technology firms, and policymakers. If those choices are made wisely, AI can support both innovation and stability.

The objective is not to hold back innovation. There are surely many benefits to be realized. And those benefits are the baseline outlook. But it is equally important to ensure that innovation takes place within a financial system that remains **trustworthy, resilient, and stable.** That requires action now, because the policy framework must evolve as quickly as the technology itself, and it can only evolve as fast as the evidence base allows. Which is, in the end, why I am speaking to this audience rather than another one.

References

Adrian, Tobias, Tamas Gaidosch, Marina Moretti, Mahvash Qureshi, and Rangachary Ravikumar. 2026. "Artificial Intelligence and Cybersecurity in the Financial Sector." IMF Note 2026/005, International Monetary Fund, Washington, DC.

AI Security Institute. 2026. "Our Evaluation of Claude Mythos Preview's Cyber Capabilities." April 13. Department for Science, Innovation and Technology, United Kingdom.

Aldasoro, Iñaki, Leonardo Gambacorta, Anton Korinek, Vatsala Shreeti, and Merlin Stein. 2024. "Intelligent Financial System: How AI Is Transforming Finance." BIS Working Paper 1194. Bank for International Settlements, Basel.

Bains, Parma, Guillermo Conde, Rangachary Ravikumar, and Emran Iskender. 2025. "AI Projects in Financial Supervisory Authorities." IMF Working Paper WP/2025/199, International Monetary Fund, Washington, DC.

Bank for International Settlements. 2024. "Artificial Intelligence and the Economy: Implications for Central Banks." Annual Economic Report 2024, Chapter III. Basel.

Bank for International Settlements. 2025. "Financial Stability Implications of Artificial Intelligence." FSI Executive Summary, June 26. Basel.

Bank of England. 2025. *Financial Stability in Focus: Artificial Intelligence in the Financial System*. London.

Basel Committee on Banking Supervision. 2021. *Principles for Operational Resilience*. Bank for International Settlements, Basel.

Brunnermeier, Markus K. 2026. "Artificial Intelligence and the Brave New World in Finance." Paper prepared for the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, August 27–29, 2026. Jackson Hole, WY.

Chang, Jennifer, and Federico Grinberg. 2024. "Advances in Artificial Intelligence: Implications for Capital Market Activities." Global Financial Stability Report, Chapter 3, October. International Monetary Fund, Washington, DC.

Crisanto, Juan Carlos, Cris Benson Leuterio, Jermy Prenio, and Jeffery Yong. 2024. "Regulating AI in the Financial Sector: Recent Developments and Main Challenges." FSI Insights 63, Bank for International Settlements, Basel.

CrowdStrike. 2026. *2026 Global Threat Report: The Evasive Adversary Wields AI*.

European Central Bank. 2024. "The Rise of Artificial Intelligence: Benefits and Risks for Financial Stability." Financial Stability Review, May. Frankfurt.

European Systemic Risk Board. 2026. *Addressing Frontier AI Models with Cyber Capabilities from a Financial Stability Perspective*. Frankfurt.

Federal Reserve Board. 2025. *Cybersecurity and Financial System Resilience Report*. Report to Congress, July. Board of Governors of the Federal Reserve System, Washington, DC.

Financial Stability Board. 2024. *The Financial Stability Implications of Artificial Intelligence*. Basel.

International Monetary Fund. 2024. "Cyber Risk: A Growing Concern for Macrofinancial Stability." Global Financial Stability Report, Chapter 3, April. Washington, DC.

International Organization of Securities Commissions. 2025. *Artificial Intelligence in Capital Markets*. Consultation Report, March. Madrid.

Shabsigh, Ghiath, and El Bachir Boukherouaa. 2023. “Generative Artificial Intelligence in Finance: Risk Considerations.” IMF Fintech Note 2023/006, International Monetary Fund, Washington, DC.