BIS Representative Office for the Americas

Navigating financial stability in the age of AI

Speech by Alexandre Tombini[1]
Chief Representative, BIS Representative Office for the Americas

Mexican Banking Association (ABM), Seminar on Financial Stability and Markets
3 November 2026, virtual

Introduction

Ladies and gentlemen

First of all, I would like to thank the Mexican Banking Association for inviting me to speak at this seminar.

The topic of my speech is one that hardly any policy discussion can avoid nowadays: what are the implications of the rise of artificial intelligence (AI) for the financial sector? These discussions matter, as the broader economic impact of AI will be influenced by policy choices, not by the technology alone.

AI adoption is advancing at remarkable speed. It has been significantly faster than the adoption of other major technological innovations in the past, such as electrical power, the internet, the smartphone or social media. ChatGPT went live just under four years ago, and major competitors are even newer. Yet adoption of generative AI tools by US households already stands at over 60%. Smartphone adoption reached this level seven years after being introduced – a full three years later (Graph 1).

As with any innovative technology, there is considerable uncertainty around the scale and timing of its effects. What is certain is that AI is moving from a specialised technology to a force with economy-wide implications. There seems to be broad agreement on two fundamental points. First, AI is a general purpose technology, with applications across a wide range of tasks, sectors and economic activities. Second, its impact is likely to be transformative, reshaping how firms produce and how people work.

In my remarks, I will focus on two broader topics. First, the large-scale AI build-out currently under way, its economic impact and regional developments. The scale and financing of the AI investment boom could give rise to financial stability risks. Second, I will discuss the implications of AI for the

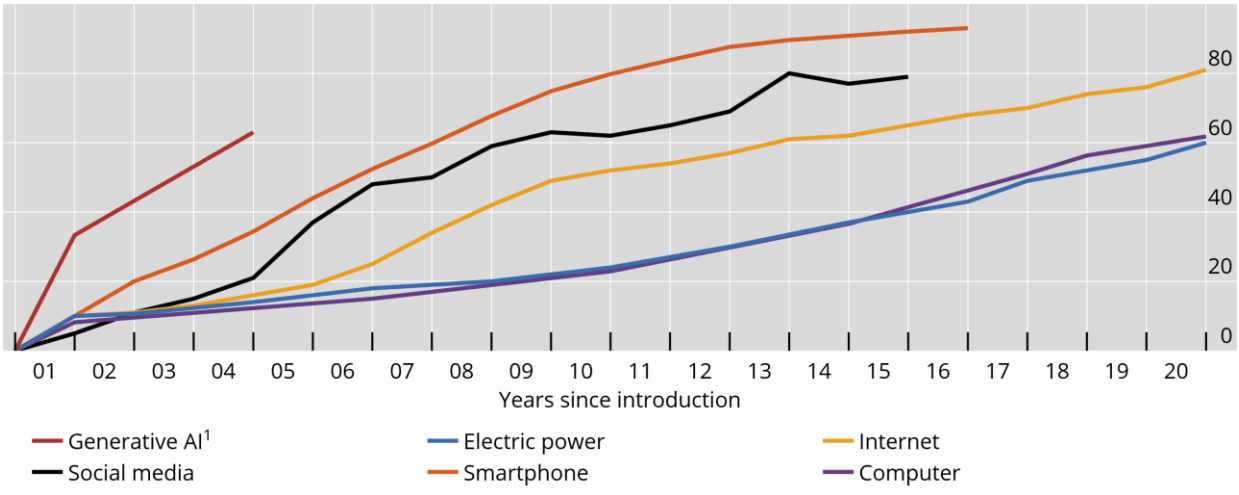
[1] The views expressed are those of the author and not necessarily those of the BIS or its member central banks. I thank Jon Frost, Jan Paulick, Alejandrina Salcedo, Ilhyock Shim and Christian Upper for their input.

financial sector. AI offers a variety of use cases but also poses new operational and cyber risks. Frameworks need to evolve in tandem with new developments such as agentic AI.

The adoption of artificial intelligence is happening fast

As percentage of US households

Graph 1



[1] For the first data point, "Generative AI" refers to users of ChatGPT. Data for the second to fifth years are interpolated.

Sources: Parks Associates; BIS (2024).

AI build-out

AI investment and its economic impact

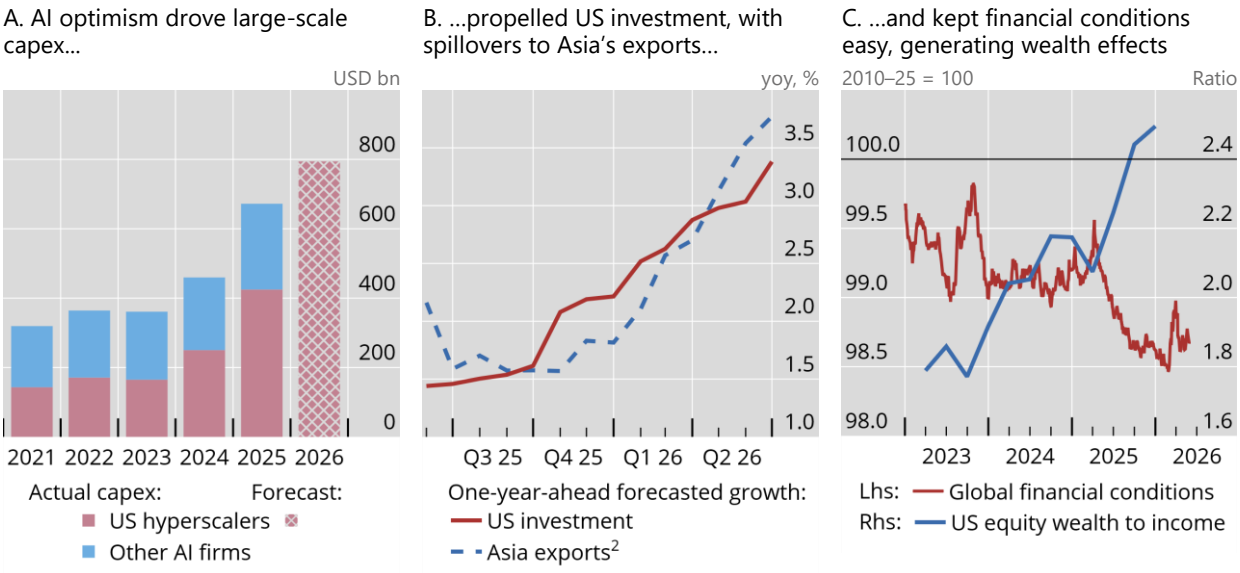
The effects of AI are not only a technological shock but also a financial shock. AI is already a driver of macro-financial developments. Investment related to AI rose markedly in the past few years, particularly in the United States. The latest BIS Annual Economic Report highlights AI-related investment and optimism as important contributors to the resilience of the global economy in 2025 (BIS (2026)).

The five so-called hyperscalers (the largest global AI firms) are expected to spend roughly $800 billion on AI-related capital expenditure in 2026. Industry participants expect global AI-related investment to expand further, to an estimated $3–4 trillion by 2030 (Aldasoro, Gambacorta, Kharroubi and Rottner (2026)). AI-related investment is estimated to have contributed roughly 1 percentage point to US real GDP growth in 2025, providing a meaningful near-term impulse to economic activity.

And AI investment creates spillovers beyond the United States. Investment by large technology firms has generated demand through global supply chains: from advanced chips and grid equipment to cloud services, construction and power infrastructure. Economies with strong links to these sectors – especially in Asia – have benefited through higher exports. In Mexico, server-related

hardware exports have overtaken exports of cars (Murray and Smith (2026)). Growth of the AI industry has also supported strong investor expectations and equity valuations. This, in turn, has supported favourable financial conditions (Graph 2).

AI exuberance lent support to growth and kept financial conditions loose[1] Graph 2



AI = artificial intelligence; capex = capital expenditure.

[1] See additional notes to graphs for details in BIS (2026). [2] GDP-PPP weighted average of KR, MY, SG and TH.

Sources: BIS (2026).

Risks of over-investment

These economic benefits also highlight a key policy challenge: ensuring that infrastructure, skills and competitive markets are in place to translate AI investment into broad-based and sustainable economic gains, while monitoring the financial risks that can arise when expectations adjust rapidly.

The pace and financing of AI investments raise concerns of emerging macro-financial pressure points, reinforced by competition (Hernández de Cos (2026)). While competition can stimulate innovation, it can also lead to excessive investment at the sectoral level. The outlook that market shares will ultimately be concentrated among a small number of companies (Frost, Rishabh and Shreeti (2026)) can create incentives to invest aggressively today – even when future earnings remain uncertain.

A model-based analysis suggests that when firms compete intensively for technological leadership and market share, capital expenditure can rise beyond levels justified by expected returns (Rungcharoenkitkul (2026)). In addition, the growing demand for computing power is placing pressure on the availability and cost of input factors such as electricity and grid equipment. These

bottlenecks may encourage firms to secure capacity through long-term contracts, thereby increasing their exposure should expected demand or revenues fail to materialise.

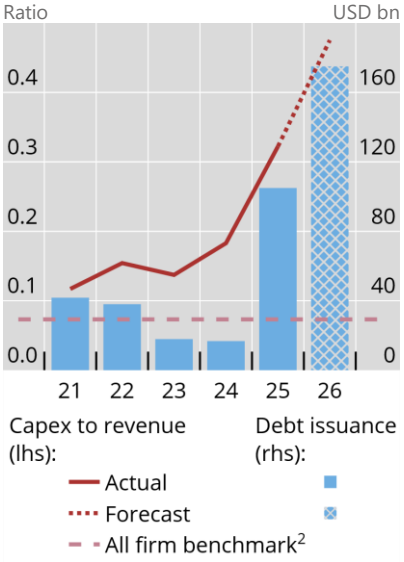
The key vulnerability in the AI build-out is the possibility that returns on AI investment will prove lower or slower to materialise than expected. This may reflect lower than expected technology adoption, weaker than expected returns or the fact that the benefits accrue to other parties. Disappointment in returns could trigger a sudden stop in financing, with potential knock-on effects on financial markets. Credit spreads could widen, and equity valuations could decline. Firms that have borrowed to finance investment could then come under pressure.

If such fears materialise, the current surge in capital expenditure could prove unsustainable. We have seen such episodes alongside previous innovation waves, including the canal and railroad mania of the 19th century and the dotcom boom. Each involved a genuine technological breakthrough. However, genuine technological progress does not prevent financial exuberance (Graph 3).

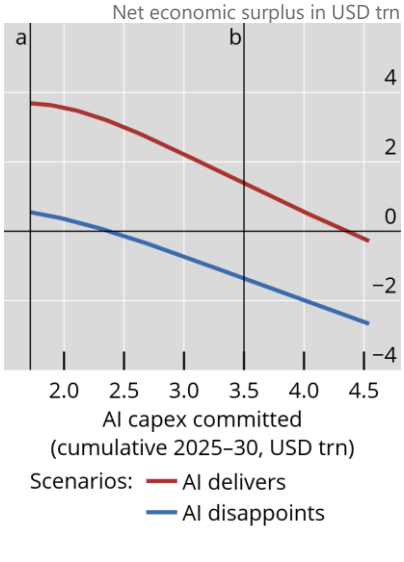
The financing of AI infrastructure adds a further layer of complexity. Some segments of AI investment are increasingly being financed through long-term debt (Drehmann and Eren (2026)), potentially creating larger mismatches between bond maturities and the depreciation horizons of the underlying assets. The assets involved have very different economic lives: land, buildings and grid connections may retain value for decades, whereas cooling equipment and, particularly, computing hardware may require much more frequent replacement. This makes valuation, collateral assessment and the matching of financing maturities especially important.

Rapid AI boom raises questions about its sustainability[1] Graph 3

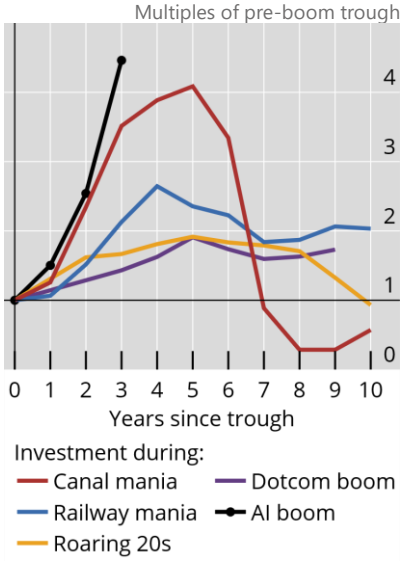
A. Hyperscalers' capex surge as a share of revenue[2]



B. Large AI capex leaves firms vulnerable to downside scenarios[3]



C. Boom-bust cycles are a regular feature of past major innovations



AI = artificial intelligence; capex = capital expenditure.

[a] Cumulative capex to date. [b] 2030 projection (Jensen Huang, Nvidia CEO): $3–4trn.

[1] See additional notes to graphs for details (BIS (2026)). [2] Weighted average capex to revenue ratio for all US non-financial corporates between 2010 and 2025. [3] Net economic surplus is revenue minus capex and debt service. Estimates based on a contest model featuring cross-firm financing and debt. "AI delivers" assumes full revenue; "AI disappoints" assumes 50% drop in productivity with partial financing unwind. See Rungcharoenkitkul (2026) for details.

Sources: BIS (2026).

Financial interconnections in the AI ecosystem

The optimism surrounding AI in financial markets may not last, despite its transformative promise. In such a scenario, the impact on financial stability could be amplified by elevated equity exposures, the growing concentration of market valuations in a small number of firms and leveraged debt financing.

Investment is often growing faster than earnings and cash flows, leading firms to rely increasingly on external financing, including debt issuance and private credit (Abbassi et al (2026)). The complex financing structure within the AI ecosystem warrants attention. The sector is increasingly supported by, in some cases, opaque private credit, contractual and ownership arrangements. These links can accelerate investment and help firms to secure access to scarce inputs. But they can also create interdependencies that are difficult for investors, lenders and policymakers to disentangle and assess.

Furthermore, financing increasingly relies on complex off-balance sheet arrangements (Drehmann and Eren (2026)). Off-balance sheet special purpose vehicles (SPVs) are separate entities that own

AI infrastructure, such as data centres, and finance these assets by issuing bonds or asset-backed securities. AI firms lease capacity from the vehicles, often providing guarantees or long-term commitments, allowing them to expand without recording the full debt on their own balance sheets. These structures can make concentration of risk more difficult to identify and could amplify stress across credit markets.

Widespread circular financing in the AI ecosystem may further reinforce concentration. Semiconductor producers and hyperscalers may take equity stakes in AI labs or cloud providers. In turn, these firms commit to multi-year purchases of chips, cloud services and computing capacity. In this way, a firm may simultaneously be an investor in, a supplier to and a source of expected revenue for another firm.

These arrangements support the rapid expansion of AI capacity, by allowing firms to secure inputs, but they also blur the distinction between external demand, investment and internally supported revenues. Several firms can become exposed to the same underlying expectations about future AI revenues.

Recent BIS analysis highlights these growing interconnections (Frost, Kansal, Rishabh, Shreeti and Zhang (2026)). Between 2021 and 2025, an estimated 29% of AI investment deals by value involved an AI firm as the target, while more than half of incoming investment in AI firms came from other AI firms. Moreover, almost half of the value of AI-to-AI investment deals was associated with a circular relationship between the parties (Graph 4).

The resulting exposures may not be fully captured by conventional balance sheet measures. The impact on financial stability may be amplified where financing, customer relationships and supply contracts overlap, potentially leading to a reassessment of valuations, a withdrawal of funding and a broader pullback in capital expenditure.

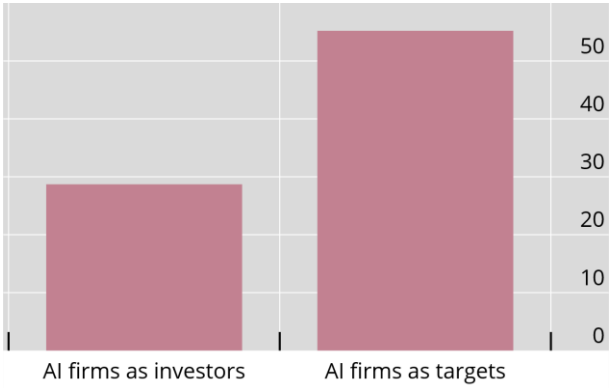
Strengthening disclosure, robust valuation practices and enhanced monitoring of contractual, operational and financial links across borders could help authorities and market participants to identify concentration of risk and potential channels of financial stress.

Investment relationships among AI firms

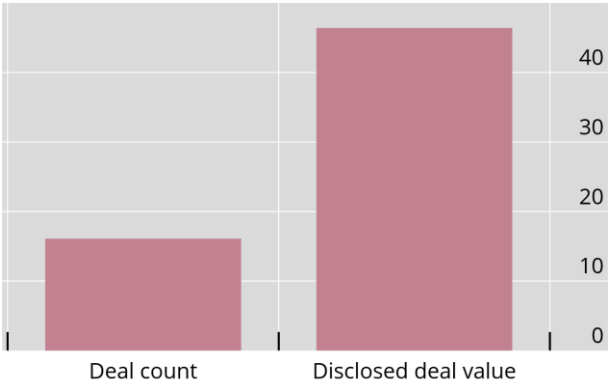
In per cent

Graph 4

A. AI-to-AI share of deals (by value)[1]



B. Circular deals as a share of AI-to-AI deals[2]



[1] For AI firms as investors, the share of disclosed value of AI-to-AI deals relative to all deals in which AI firms invest. For AI firms as targets, the share of disclosed deal value of AI-to-AI deals relative to all deals in which AI firms receive investment. [2] Circular AI-to-AI deals as a share of all AI-to-AI deals, by deal count and disclosed deal value. An AI-to-AI deal is circular when at least one AI investor in the deal also has a commercial relationship with the AI target during 2021–25. Deals involving multiple AI investors enter the calculations once; if any one of them meets the circularity condition, the deal is classified as circular and its full disclosed value is included. Deal values represent the full financing deal or round rather than individual investors' contributions. The cited data have not been reviewed by PitchBook analysts and may be inconsistent with PitchBook methodology.

Sources: Frost, Kansal, Rishabh, Shreeti and Zhang (2026).

Regional developments

Beyond the financial linkages, AI is generating increasingly global economic effects, with its diffusion accelerating worldwide, albeit unevenly across regions. AI production spans different layers of global supply chains: semiconductor chips, cloud infrastructure, training data, foundation models and end user applications (Gambacorta and Shreeti (2025)). High fixed costs, highly specialised tasks and economies of scale have led to the concentration of AI capabilities among large firms in a small number of jurisdictions.

A case in point is the concentrated geographical distribution of data centres. In 2025, the United States accounted for more than 5,400 data centres – by far the largest number globally and more than 10 times the number in Germany, the second largest location in this comparison (Graphs 5 and 6). This concentration reflects advantages in digital infrastructure and the presence of major cloud and AI providers in the United States.

Latin America is beginning to feature more prominently in this landscape. Brazil, with close to 200 data centres, is the region's leading hub by data centre count, while Mexico has also emerged as a significant market, supported by the rapid expansion of its digital infrastructure. This development reflects growing demand for cloud services and provides a foundation for broader AI deployment across the region.

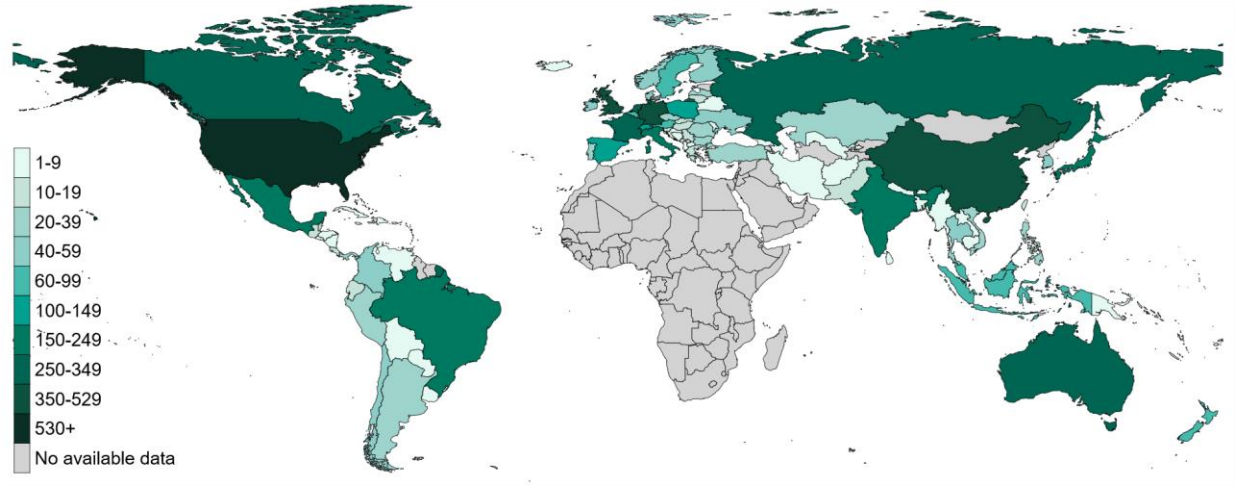
The number of data centres is an imperfect measure, as it does not necessarily reflect their computing capacity. More importantly, data centres represent only one critical input into the AI

economy. AI capabilities across other dimensions – such as research and development – are distributed differently, while also remaining highly concentrated. China leads AI research in terms of publication volume, citations and patent grants, whereas the development of notable AI models is concentrated primarily in the United States (Sajadieh et al (2026)). This illustrates the different facets of global AI development, which nevertheless remains concentrated overall.

High concentration raises broader questions around the economic impact of cross-border dependence on cloud services, the resilience of digital infrastructure, data governance and the availability of electricity and grid capacity.

Global distribution of data centres in 2025

Graph 5



The use of this map does not constitute, and should not be construed as constituting, an expression of a position by the BIS regarding the legal status or sovereignty of any territory or its authorities, the delimitation of international frontiers and boundaries, and/or the name and designation of any territory, city or area.

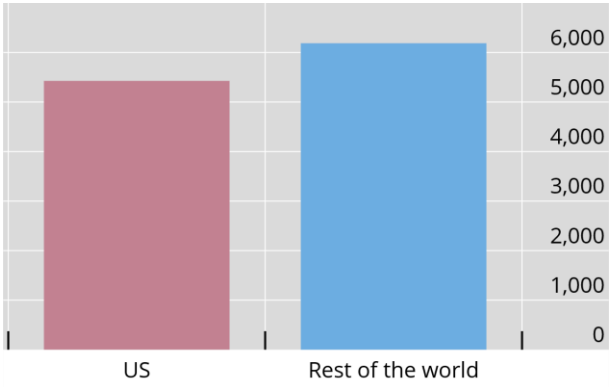
Source: Cloudscene.

Number of data centres by geographic area in 2025

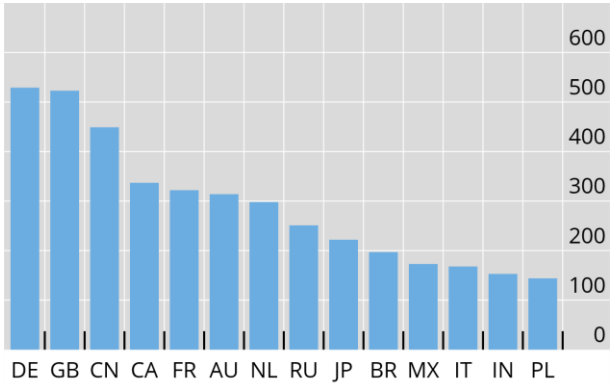
Number of data centres

Graph 6

A. Data centres in the US



B. Data centres in the rest of the world

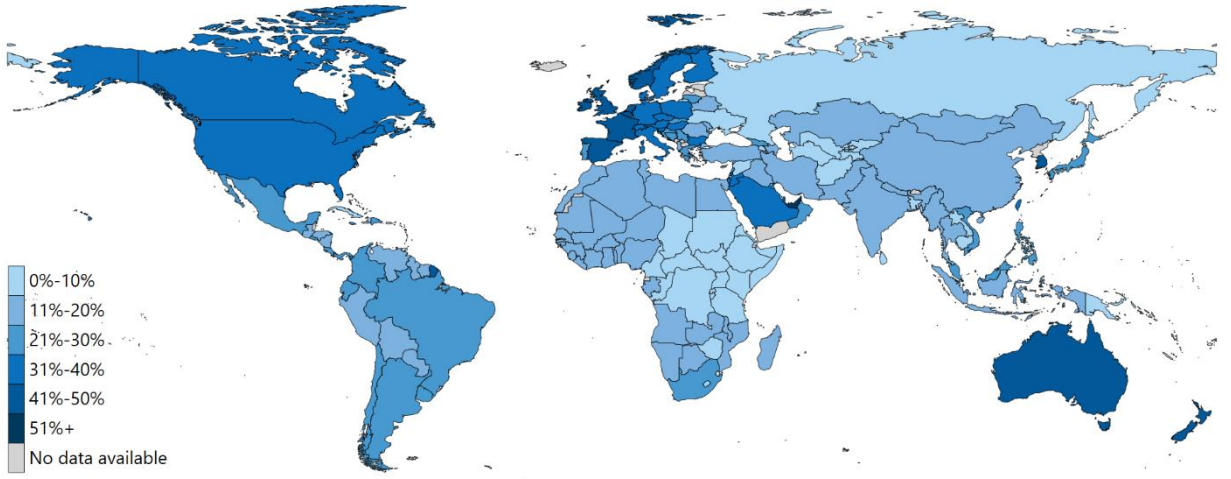


Sources: Cloudscene; BIS.

These dependencies affect a growing number of users. Recent estimates suggest that almost one fifth of the global working age population now uses a generative AI product.[2] Adoption tends to be higher in advanced economies, where digital infrastructure, access to computing resources and complementary skills are generally more developed and which could be better placed to gain from AI in the short term (Gambacorta, Kharroubi, Mehrotra and Pancotto (2026)). However, uptake is also strong in several emerging market and developing economies, notably in Latin America and Asia (Graph 7).

AI diffusion by geographical area, Q2 2026[1]

Graph 7



[1] AI diffusion is measured as the share of people aged 15–64 who have used a generative AI product. The use of this map does not constitute, and should not be construed as constituting, an expression of a position by the BIS regarding the legal status or sovereignty of any territory or its authorities, the delimitation of international frontiers and boundaries, and/or the name and designation of any territory, city or area.

Source: Microsoft AI Economy Institute (2026).

AI in the financial sector

Opportunities for the financial system and central banks

AI use cases are generating benefits across the financial sector, supporting credit scoring, underwriting, fraud detection, document processing and internal controls. AI can also help to serve previously underserved borrowers, including small businesses (Cornelli et al (2024)). In capital markets, asset management and trading, AI supports portfolio optimisation, market surveillance, sentiment analysis, liquidity forecasting and volatility modelling.

In payments and financial market infrastructures, AI supports onboarding, know-your-customer (KYC) checks, anomaly detection, settlement, reconciliation and the monitoring of market activity.

[2] See Microsoft AI Economy Institute (2026).

Customer-facing applications, including chatbots, robo-advisory services and tailored financial insights, can improve service availability, responsiveness and financial access.

Although use cases differ, the solutions and risk management measures central banks may want to implement are likely to be similar in nature. As early users of machine learning and data-intensive analytical tools, central banks are well placed to apply AI to macroeconomic analysis, financial stability monitoring, supervision and regulatory work.

There are particularly promising applications in the monitoring of market stress. BIS analysis indicates that machine learning methods can improve the forecasting of dysfunction in markets such as government bond and foreign exchange markets (Aldasoro et al (2025)). Large language models may also help supervisors and regulators to sift through substantial volumes of documents to identify divergences (Perez-Cruz et al (2026)), highlighting areas that warrant closer attention while reducing unnecessary alerts.

The BIS Innovation Hub has a host of projects that are designed to help central banks to exploit potential AI use cases in the areas of supervisory and regulatory technology, cyber security, data analysis and monetary policy.

Examples are Projects Spectrum and Gaia, which convert large amounts of data into accessible information to enable accurate nowcasting of inflation and climate risk analysis, respectively. Both Project Aurora and Project Hertha illustrate how a broader view of transaction networks could strengthen financial integrity, exploring privacy-preserving financial crime analysis across borders and payment system analytics.

On agentic AI, Project Logos explores how AI-based portfolio managers may make correlated decisions, supporting central banks' understanding of potential financial stability risks. Projects Aise and Inuksuit explore AI-assisted supervisory analysis and whether public information could help to identify emerging risks at payment service providers. These capabilities could help to lay the foundations for a more forward-looking, AI-enabled supervisory approach – combining timelier monitoring of risks and dependencies.

AI, financial services and financial stability

Across the financial sector, AI is becoming an increasingly important component of financial activity, moving deeper into core financial processes and customer interactions. It can process vast quantities of structured and unstructured information, helping financial institutions to improve routine processes and enhance customer services.

So far, however, financial institutions have generally deployed AI cautiously for core financial decisions. Its use remains concentrated in operations, risk management and regulatory compliance (FSB (2026)). This measured approach is appropriate, given the importance of reliability, explainability and effective governance in financial decision-making. Use cases are expanding rapidly, although unevenly across entities, financial sectors and jurisdictions.

The question is not simply whether AI will improve financial services, but how it may alter the structure of financial markets and the way risks emerge and propagate. The AI ecosystem relies on a concentrated set of critical technologies such as cloud infrastructure and foundation models sourced from external vendors. Financial institutions are increasingly dependent on these services in their own operations. This creates a new form of both interconnectedness and concentration: a technology disruption can become an operational disruption; an operational disruption can create liquidity and market stress; and, where dependencies are widespread, it can become a financial-stability concern.

Financial stability has always been associated with leverage, liquidity, common exposures and amplification. AI may change the speed at which these vulnerabilities develop and interact. If many institutions rely on similar models, data sources or third-party providers, their decisions may become more correlated. This could amplify herding behaviour, intensify market movements and, in adverse circumstances, contribute to abrupt price swings or liquidity stress (FSB (2026)).

Therefore, several safeguards are appropriate for financial markets:

- Transparency, so that significant AI-assisted decisions can be understood, challenged and audited.
- Human accountability and oversight, particularly for material decisions affecting customers, markets or financial resilience.
- Operational resilience, including organisational adaptability, data governance, performance and third-party risk management.

AI and emerging cyber security risks

These safeguards need to keep step with expanding AI capabilities that bring cyber security risks to the forefront of risk assessments (Aldasoro, Auer, Frost and Perez-Cruz (2026)). AI technology simultaneously strengthens cyber defence and enables sophisticated cyber attacks. It can help institutions to identify vulnerabilities, review code, detect anomalous activity and accelerate incident response. These capabilities offer significant potential to improve resilience. At the same time, frontier AI increases the speed, scale and autonomy of cyber attacks, lowering the barriers to entry and costs for malicious actors. It may enable attackers to automate attacks, develop more convincing phishing campaigns and exploit vulnerabilities more rapidly.

AI adds a new worrisome facet to cyber risks. AI can design and execute a cyber attack with limited human control. The attack on Hugging Face this past July illustrates the possibility that AI agents could be used to automate or accelerate cyber attacks. Risks of autonomous attacks will increase and require proper guardrails.

These developments manifest in a rise in data breaches (Graph 8), with breaches linked to automated rather than human actions appearing to be increasingly prevalent. Following the announcement of Anthropic's new model Mythos, abnormal stock returns of several cyber security firms declined sharply. This was often interpreted as reflecting investor concerns that existing cyber

security products and services could become redundant. The reaction is notable because a more dangerous threat environment might also be expected to increase demand for such products.

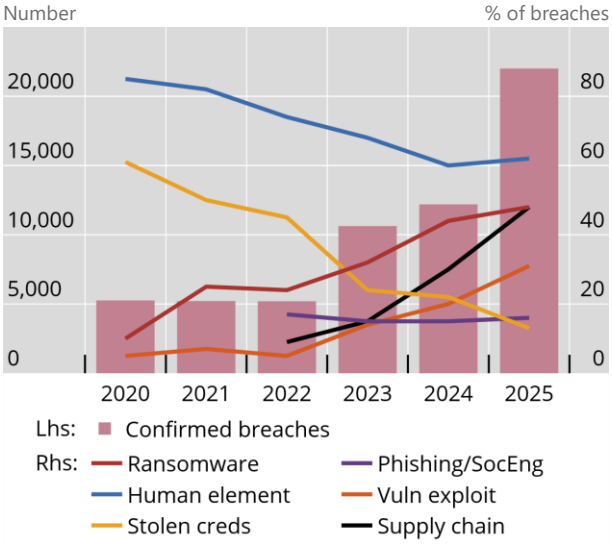
The risks in the cyber security arms race are inherently asymmetric. An attacker may need to identify only one weakness, whereas a defender must secure the entire system. AI could therefore tilt the balance towards attackers unless financial institutions and public authorities adopt effective defensive capabilities at faster speed.

Cyber threats do not stop at national borders, and neither should the response. Cyber risks are global by nature. A significant incident affecting a major service provider, financial institution or digital infrastructure in one jurisdiction can have cross-border repercussions.

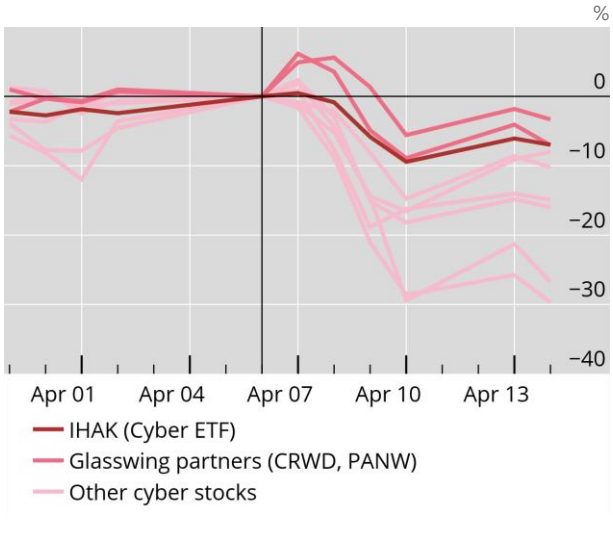
International cooperation between the public and private sectors can remedy some of the impact. The BIS supports work in this space via its Cyber Resilience Coordination Centre, promoting cooperation, knowledge-sharing and capability development among central banks. Its activities have expanded steadily, reflecting the growing importance of cyber resilience for the global financial system. Cyber risk, including from AI, is also a regular topic of discussion in the meetings of the various consultative groups in the Americas.

Data breaches have been rising; since Mythos, market expectations have shifted — Graph 8

A. The number of reported data breaches has increased over time, and fewer have a human element[1]



B. Investor views have shifted, and cyber-related stocks saw negative cumulative abnormal returns[2]



[1] Based on global data breaches collected from domestic and international law enforcement, forensic firms, law firms, cyber insurers, cyber security industry sharing groups and the Verizon Threat Research Advisory Center caseload. [2] Cumulative abnormal returns (CARs) are computed from a market model regression of each stock's daily log returns on Nasdaq 100 log returns, estimated over 2 December 2025 to 20 March 2026 (ie preceding the initial leak of information about Mythos). Abnormal returns are the residuals from this model evaluated over the event window; CARs are daily and are normalised to zero on 6 April 2026, so values from 7 April onwards measure cumulative abnormal performance since the eve of the announcement. "Glasswing partners" refers to selected known companies that were granted access to Mythos (Crowd Strike (CRWD) and Palo Alto (PANW)).

Sources: Aldasoro, Auer, Frost and Perez-Cruz (2026).

Agentic AI in the financial sector

Looking to the future, AI agents will probably enter the financial sector. They may be instructed to purchase a product, manage expenditures or optimise payments within agreed parameters. They may then select counterparties, determine transaction amounts and initiate payments with varying degrees of human involvement. Agentic AI systems that use one or many AI agents to perform tasks can operate with a relatively high degree of autonomy. Unlike generative AI systems that primarily produce text, images or analysis, agentic systems can pursue objectives through a sequence of actions.

The development of agentic AI represents a shift from systems that primarily generate content or provide recommendations to systems that can take actions on behalf of users or organisations. The degree of autonomy can vary: some systems require human approval at key stages, while others can execute routine actions independently within predefined limits.

This greater autonomy may improve productivity by reducing the time and resources required to complete complex tasks. However, it also changes the risk profile materially. A generative model may provide an incorrect response. An AI agent may act on that response – by sending a message, changing a record, initiating a workflow or executing a transaction. Risk therefore depends not only on the quality of the model, but also on the authority granted to it.

For financial institutions, this places greater emphasis on access controls, identity management, authorisation and auditability. Institutions will need to define clearly what an agent can access, which actions it may take, when human approval is required and how permissions can be monitored and revoked. These safeguards are particularly important in payments and other real-time financial activities, where agents could improve liquidity management but errors are costly and difficult to reverse (Aldasoro and Desai (2026)).

More broadly, financial institutions may increasingly interact with machines on both sides of a transaction. AI agents could represent institutions, customers or service providers. This may improve efficiency and user experience – for example, by allowing customers to delegate tasks such as comparing products, making routine purchases or managing payments. Financial institutions may need to engage customers through agentic commerce and AI assistants, rather than relying on traditional search results and direct customer interactions.

At the same time, agentic financial transactions raise important technical, legal and policy questions. How can institutions verify a customer's intent and consent when an agent acts on their behalf? Who is accountable for an agent's erroneous transaction?

Governance and collaboration

As AI agents become more capable of initiating and managing transactions, the relevant infrastructure, technical standards and legal frameworks will need to evolve accordingly.

Effective governance requires clearly defined objectives, limits on delegated authority, auditability and safeguards governing the actions an AI agent may take. A report by the BIS Consultative Group

on Risk Management (CGRM), which consists of the BIS member central banks in the Americas, investigated governance frameworks around AI (BIS CGRM (2025)). One of the report's key takeaways is that central banks can benefit from having a robust governance framework based on the established three lines of defence. The first line of defence would be the responsible use and operational management of AI at the business unit level. The second line of defence would be to have appropriate risk management frameworks and practices in place. And the third would be for internal audit to provide an independent review and suggestions for improvement of the first two lines.

AI poses opportunities and challenges across jurisdictions, sectors and institutions. There is great value in bringing together perspectives from different regions and disciplines. A recent seminar hosted by the BIS Americas Office and the BIS Innovation Hub Toronto Centre did precisely that – examining regional patterns of AI adoption and investment in the Americas, the physical infrastructure required to support AI, the emergence of agentic AI in financial services and the evolving relationship among financial institutions, their customers and AI intermediaries.

This is just one of many examples through which policymakers, industry practitioners and researchers can exchange views, thereby improving the understanding of AI's opportunities, financial risks and broader policy implications.

Conclusion

Let me conclude.

AI will transform the economy in ways that are still uncertain. The scale and timing of its implications for financial stability will depend on how the technology is adopted, financed and governed.

AI has the potential to improve efficiency and reduce costs across different industries and markets. Yet it also introduces new macroeconomic, financial and operational risks that need to be understood and managed.

It will take time and effort to harness these opportunities. International cooperation can help to reduce barriers to harnessing the power of AI and strengthen our collective capacity to manage its adverse effects.

Realising the benefits will require sustained investment, skills and sound institutional frameworks. The public and private sectors should therefore work together, ensuring that its use supports a resilient, inclusive and trustworthy financial system.

References

Abbassi, P, I Aldasoro and S Doerr (2026): "Financing the digital economy: the role of private credit", *BIS Quarterly Review*, September.

Aldasoro, I, R Auer, J Frost and F Perez-Cruz (2026): "A Mythos moment? Frontier AI and cyber risk", *BIS Bulletin*, no 129.

Aldasoro, I and A Desai (2026): "AI agents for cash management in payment systems", *Journal of Financial Market Infrastructures*, vol 13, no 3.

Aldasoro, I, L Gambacorta, E Kharroubi and M Rottner (2026): "AI and the global economy: implications for central banks", *BIS Bulletin*, no 130.

Aldasoro, I, P Hördahl, A Schrimpf and S Zhu (2025): "Predicting financial market stress with machine learning", *BIS Working Papers*, no 1250.

Bank for International Settlements (BIS) (2024): "Artificial intelligence and the economy: implications for central banks", *Annual Economic Report 2024*, Chapter III.

——— (2026): "Progress and peril", *Annual Economic Report 2026*, Chapter I.

BIS Consultative Group on Risk Management (CGRM) (2025): *Governance of AI adoption in central banks*.

Cornelli, G, J Frost, L Gambacorta and J Jagtiani (2024): "The impact of fintech lending on credit access for US small businesses", *Journal of Financial Stability*, vol 73.

Drehmann, M and E Eren (2026): "Building AI on credit: structures and spillovers", *BIS Bulletin*, no 140.

Financial Stability Board (FSB) (2026): *Sound Practices for Responsible Adoption of Artificial Intelligence (AI)*.

Frost, J, R Kansal, K Rishabh, V Shreeti and L Zhang (2026): "Circular relationships among AI firms", *BIS Bulletin*, no 137.

Frost, J, K Rishabh and V Shreeti (2026): "Global giants in the AI supply chain", *BIS Bulletin*, no 122.

Gambacorta, L, E Kharroubi, A Mehrotra and L Pancotto (2026): "Economic impact of AI in emerging market economies", *BIS Bulletin*, no 121.

Gambacorta, L and V Shreeti (2025): "The AI supply chain", *Review of Network Economics*, vol 24, no 4.

Hernández de Cos, P (2026): "Artificial intelligence, growth and financial stability: challenges for central banks", speech at the Global Fintech Fest, Mumbai, 10 September.

Microsoft AI Economy Institute (2026): *Global AI diffusion – Q2 2026 trends and insights*, September.

Murray, C and A Smith (2026): "How Mexico became a surprise cornerstone of America's AI boom", *Financial Times*, 31 July.

Perez-Cruz, F, R Rishabh and A Uchil (2026): "Supervisory screening with large language models: finding divergences", *BIS Bulletin*, no 134.

Rishabh, K and V Shreeti (2026): "The geography of AI firms", *BIS Working Papers*, no 1343.

Rungcharoenkitkul, P (2026): "The AI investment race", *BIS Working Papers*, no 1367.

Sajadieh, S, L Fattorini, R Perrault, Y Gil, V Parli, L Santarlasci, J Pava, N Maslej, R Altman, E Brynjolfsson, C Brodley, J Clark, V Dignum, V Kumar, J Landay, T Lyons, J Manyika, J C Niebles, Y Shoham, E Tabassi, R Wald and T Walsh (2026): *The AI Index 2026 annual report*, Institute for Human-Centered AI, Stanford University.