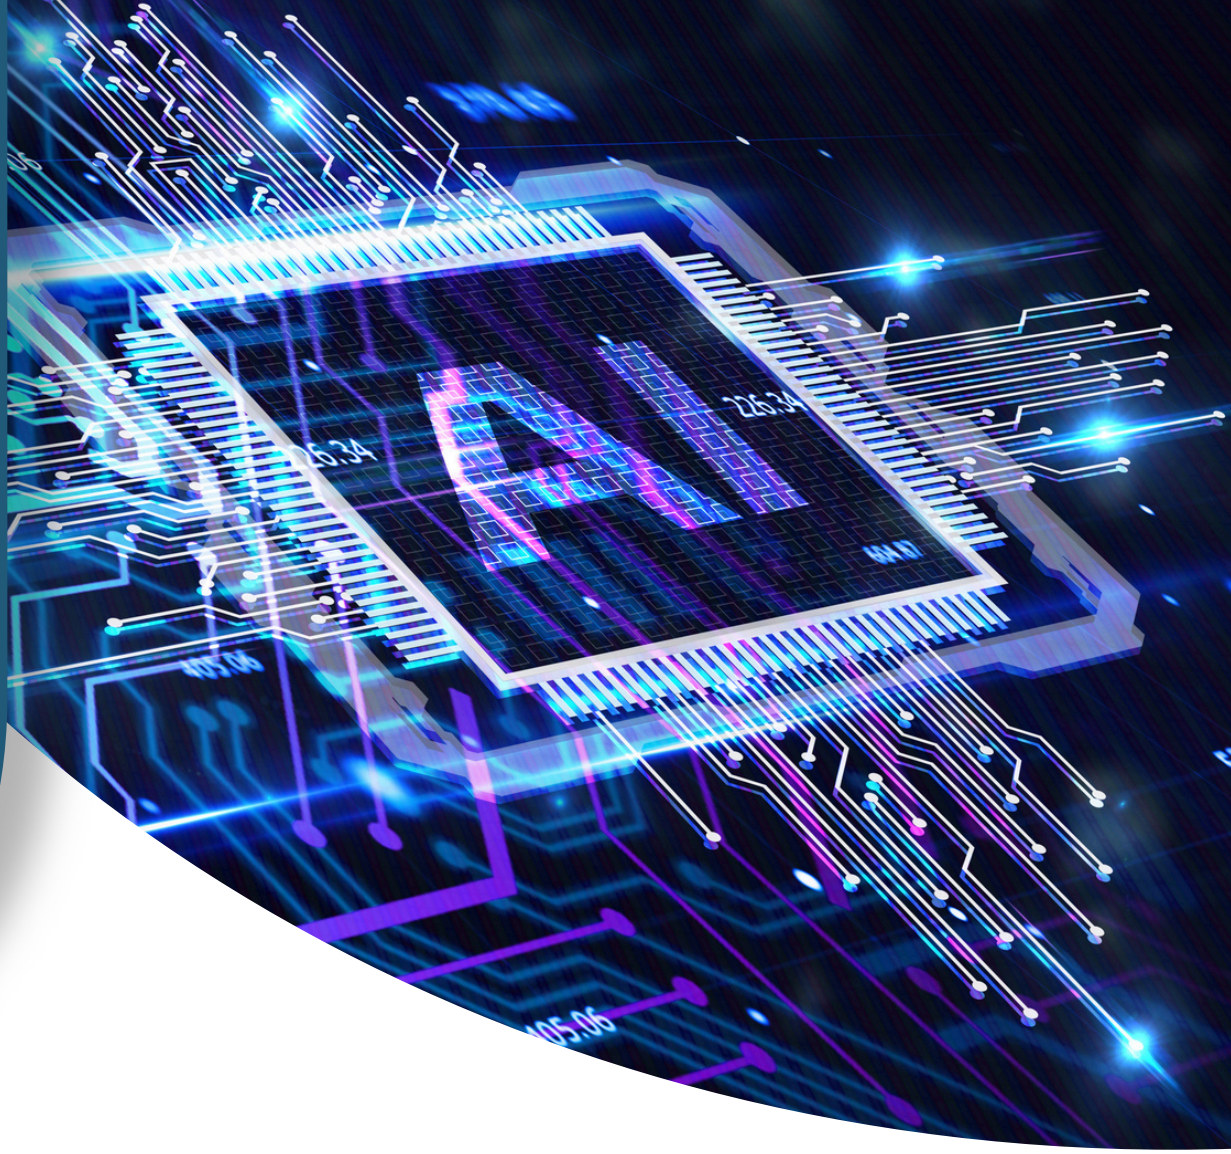


ECMI IN-DEPTH ANALYSIS

THE FINANCIAL FAULT LINES OF THE AI BOOM: FROM INTERCONNECTED FINANCING TO SYSTEMIC RISK

Judith Arnal and Apostolos Thomadakis

ECMI EUROPEAN CAPITAL MARKETS INSTITUTE

November 2026 - 02

SUMMARY

The AI investment boom has been a boon for financial markets, sending technology stocks to new highs and fuelling investment in data centres and computing infrastructure. But the financial risks extend beyond technology stocks: a correction in AI-related valuations could also affect credit markets and financial institutions.

There are three potential vulnerabilities in particular: concentration and interdependence along the AI value chain; circular financing between suppliers and customers; and the growing reliance on debt, long-dated leases, special purpose vehicles and private credit. Such vulnerabilities may amplify financial linkages through which a correction in expectations about AI investments could transmit credit losses, refinancing problems, diminishing collateral values and liquidity strains, with potential spillovers beyond the AI sector.

Europe has risks in the form of holdings of US equities and corporate debt, participation in private financing structures and links to the AI supply chain. The available evidence does not allow a systemic risk to be ruled out: AI-related exposures are large and growing, and their size, location and interconnectedness are difficult to assess.

The response should start at the international level, because no single jurisdiction can close the data gaps on its own. The Financial Stability Board should promote comparable disclosures, cross-border sharing of supervisory data and coordinated scenario analysis. European authorities should then make better use of existing supervisory data to identify interlinked AI-related exposures and integrate AI shocks into their financial stability analyses. Higher transparency and more knowledge on loss propagation would contribute to stronger financial resilience, without the need for a specific regulatory regime for AI financing.

The European Capital Markets Institute (ECMI) is an independent, non-profit research institute that conducts in-depth research aimed at informing the public debate and policymaking process on a broad range of issues related to capital markets. ECMI's operations and staff are managed by the Centre for European Policy Studies (CEPS).

Judith Arnal is Associate Senior Research Fellow at CEPS and ECRI (the European Credit Research Institute). Apostolos Thomadakis is Head of Research at ECMI, and Senior Research Fellow and Head of the Financial Markets and Institutions Unit at CEPS.

© Copyright 2026, Judith Arnal and Apostolos Thomadakis. All rights reserved.

INTRODUCTION

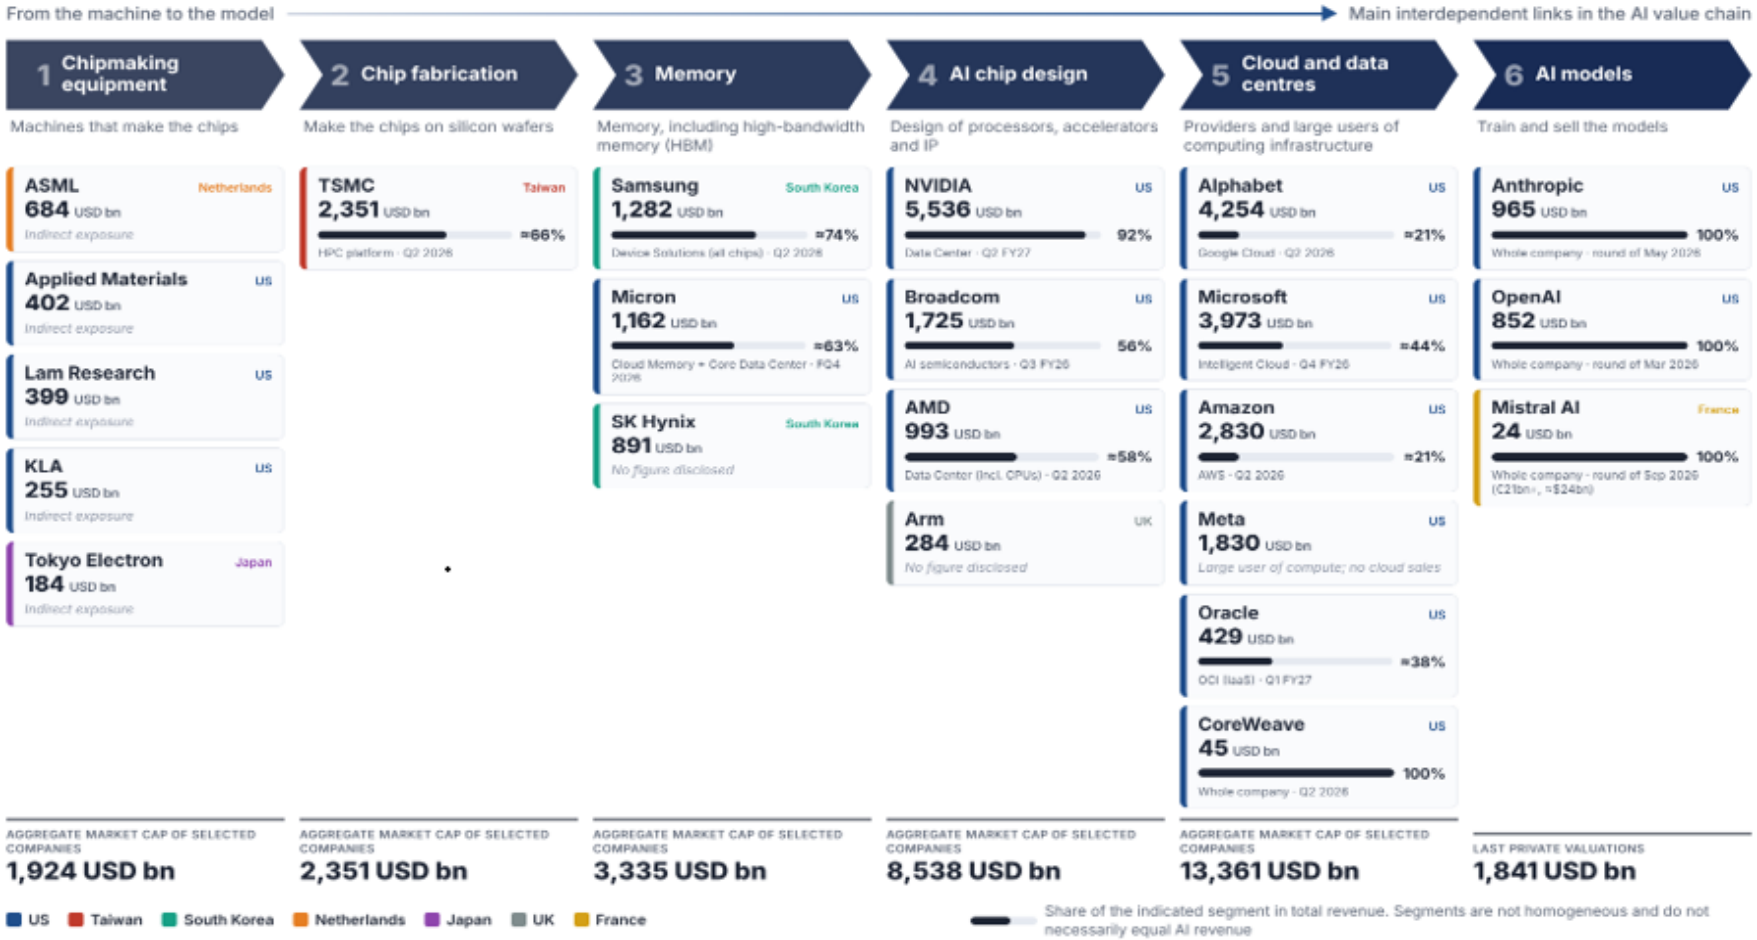
The artificial intelligence (AI) boom is transforming global financial markets, with the largest technology companies enjoying record-breaking valuations, and investments in AI infrastructure growing exponentially. NVIDIA alone has a market capitalisation of approximately USD 5.5 trillion, and the combined valuation of the major listed companies in the AI value chain comes close to USD 30 trillion (see Figure 1). Amazon, Alphabet, Microsoft and Meta expect to spend approximately USD 720-745 billion on capital expenditures in 2026, while J.P. Morgan estimates cumulative AI investment to total USD 5.5 trillion by 2030, of which USD 4.1 trillion will be funded by debt.

These figures are spurring discussion about the sustainability of AI valuations and expectations of future investment. However, the impact on financial stability is more than just the risk of a correction in stock prices. An equally important question is who ultimately bears the financial risks and how losses could spread through the financial system. The International Monetary Fund (IMF) refers to the prospect of an AI-driven market correction (if expectations for investment or profits are disappointed), as a source of financial instability (IMF, 2026). How the AI investment boom is being funded, what the extent of financial exposures and their interconnections are, and what would happen if the anticipated returns on investment were not realised are growing concerns.

The financing of AI investment is changing. While the largest technology companies remain highly profitable and fund a significant proportion of their investment from internally generated cash flows, the use of debt financing is becoming increasingly important. As well as conventional corporate bonds, companies are extensively using long-term leases, special purpose vehicles (SPVs) and private credit to fund data centres, computing capacity and associated infrastructure. The Bank of England identifies the rapid growth in credit financing for AI infrastructure, including the increasing use of private markets, securitisation and off-balance-sheet financing arrangements (Bank of England, 2026). At the same time, financial relationships within the AI value chain are becoming more sophisticated. Major suppliers are investing in AI developers that, in turn, commit to purchase computing capacity, chips and other services from those suppliers. These arrangements add financial links to existing commercial relationships.

This matters because the consequences of a correction depend partly on how the boom is financed. It has been shown that asset price booms that coincide with rapid credit expansion are associated with higher financial crisis risks and more severe subsequent downturns (Jordà et al., 2015). A correction in the equity market would predominantly result in valuation and wealth losses. However, if the investments are undertaken on the basis of significant amounts of debt and integrated financing arrangements, losses on the expectation of revenues could also give rise to credit losses, difficulties in refinancing, deteriorating collateral values, and forced asset sales. These impacts could have spillover effects to other companies, particularly if there are concentrated, leveraged, or opaque exposures.

Figure 1: Main interdependent links in the AI value chain



Notes: Market capitalisation of listed companies at close of 9 October 2026 (USD billion). Private company valuations (OpenAI, Anthropic and Mistral AI) refer to their latest funding rounds. Revenue exposure percentages are based on company filings for the latest reported periods and refer to different business segments; they are not directly comparable across companies.
Source: Authors' elaboration based on CompaniesMarketCap (market capitalisations, including USD conversions for non-US listings), company filings (revenue exposure) and reported funding rounds (private-company valuations).

These risks could become systemic and pose a growing build-up of exposures. AI-related exposures are large and growing, with a rising share financed with debt and off-balance-sheet vehicles, and supervisors are unable to measure them. The big AI companies are in good health financially, but this does not protect the lenders and investors in such vehicles. The question is whether a large enough correction could reverberate through the financial system, with impacts on financial institutions and the availability of financing beyond the AI industry. Taking into account this possibility involves going beyond the listed companies and their balance sheets to the funds, insurers, pension funds, banks and financing structures that ultimately own or finance AI-related exposure. The Financial Stability Board (FSB) has identified the deepening connections between private credit funds, banks and insurers, while also pointing to important gaps in supervisory information (FSB, 2026).

Europe is directly exposed to the risks, even though big AI companies are primarily American. European investors own a substantial share of US equities and corporate bonds, and insurers, pension funds and investment funds also account for private credit and other financing structures. The European Central Bank highlights significant euro area investor exposures to highly valued US technology and AI-related equities (ECB, 2026), while its analysis documents the rise of US hyperscalers in the euro area corporate bond markets (Duquerroy et al., 2026). European companies play a vital role in the AI supply chain as well. Hence, the correction could impact the European financial markets through various channels including portfolio valuations losses, credit exposures and market adjustments. However, it is difficult to assess the size and structure of AI-specific exposures, in particular where financing takes place through multiple lenders, jurisdictions and off-balance-sheet structures.

Against this backdrop, this paper seeks to explore how the funding of the AI boom is changing, identify the vulnerabilities that may arise as a consequence of financial concentration, circular financing and growing indebtedness, and trace the channels through which a correction could propagate through the financial system. It then looks at the implications for Europe and the policy priorities for enhancing transparency, monitoring interlinking exposures and building financial resilience, both at the level of the EU and through international cooperation.

HOW THE AI BOOM IS CREATING NEW FINANCIAL VULNERABILITIES

Three aspects of the AI investment boom should be singled out: the concentration of the value-added chain of AI, the financial circularity of suppliers and customers, and the increased use of debt and complex financing structures. The following section focuses on the degree of these tendencies and their interplay.

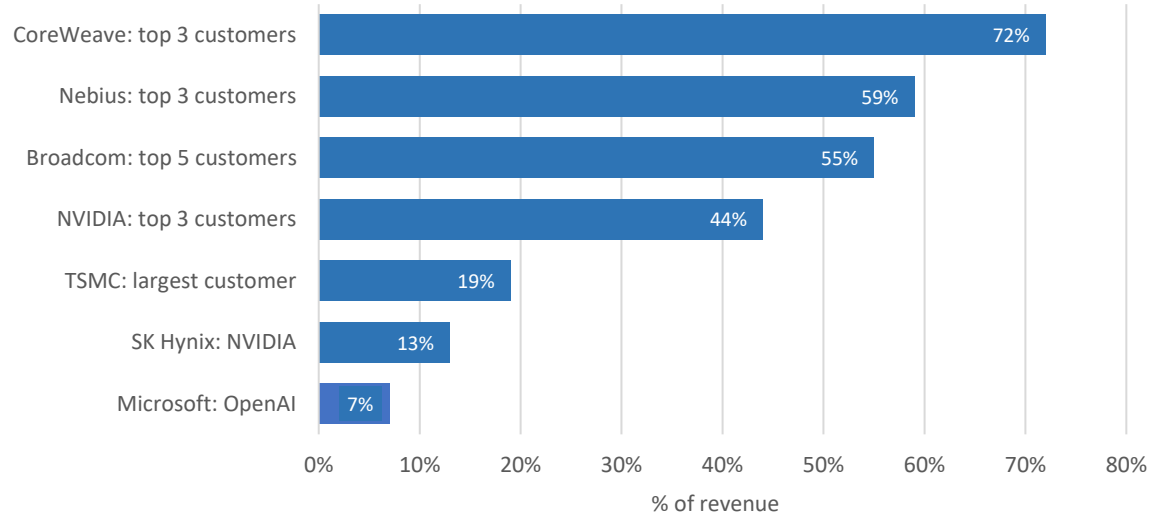
CONCENTRATION AND INTERDEPENDENCE IN THE AI VALUE CHAIN

The AI value chain includes a small set of highly integrated companies, ranging from semiconductor equipment manufacturers and chipmakers to cloud infrastructure providers and AI model developers (OECD, 2026). Equipment manufacturers such as ASML and Tokyo Electron supply the machines needed to produce advanced semiconductors. TSMC makes

accelerators designed by NVIDIA, AMD and Broadcom, while Samsung, SK Hynix and Micron supply high-bandwidth memory. Cloud providers, including Amazon, Microsoft, Alphabet and Oracle, buy computers and provide the processing power needed to train and run AI models.

This structure creates substantial commercial interdependencies. Many companies derive a large share of their revenues from AI-related activities, while depending on a limited number of major customers. NVIDIA's data-centre business accounts for 92 % of its revenues, while high-performance computing represents 66 % of TSMC's revenues and AI semiconductors 56 % of Broadcom's. Customer concentration is also considerable (see Figure 2). For example, CoreWeave's three largest customers accounted for 72 % of its revenues in the second quarter of 2026, while other major AI-linked companies also depend on a relatively small number of customers.

Figure 2: Customer concentration in the AI value chain (% of revenue)

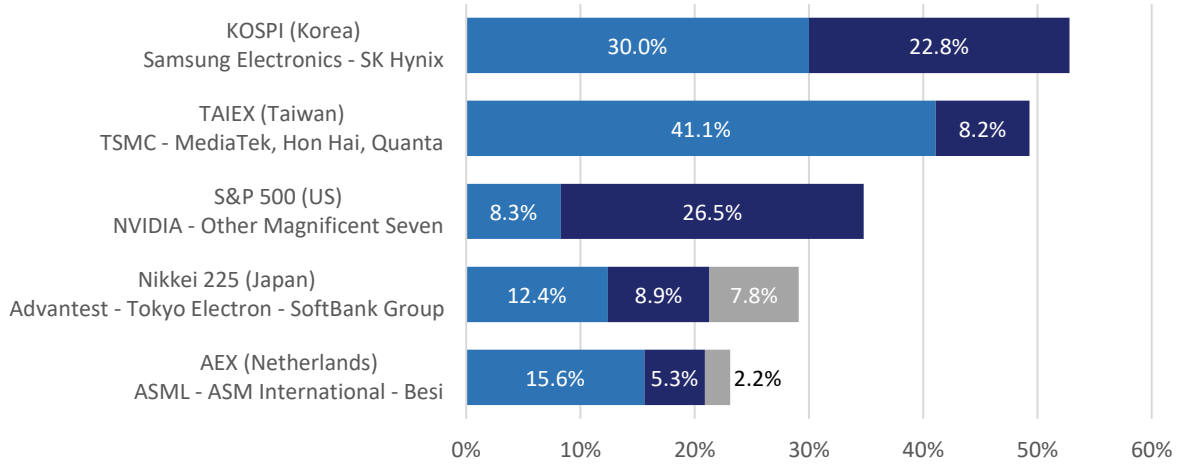


Notes: Reporting periods: CoreWeave and Nebius (Q2 2026); Broadcom (Q3 FY2026); NVIDIA (H1 FY2027); TSMC (2025); SK Hynix (H1 2026); Microsoft (FY2026). Figures refer to the share of revenue attributable to the indicated customers. Measures differ across companies in the number and identification of customers and whether they refer to direct or end customers. Comparisons should therefore be interpreted with caution.
Source: Authors' elaboration based on company regulatory filings (10-K, 10-Q and annual reports) and financial reporting.

These dependencies mean that a slowdown in AI investment would not affect companies independently. If AI developers reduced their demand for computing capacity, cloud providers could scale back their infrastructure investment, reducing orders for accelerators, memory and semiconductor manufacturing equipment. The impact could therefore propagate throughout the value chain, particularly where companies depend heavily on AI-related revenues or a limited number of customers.

This concentration is also reflected in equity markets (see Figure 3). The seven largest US technology companies make up almost 35 % of the S&P 500, with NVIDIA alone accounting for 8 %. Similar concentrations exist in other markets, mainly in South Korea and Taiwan, while in Europe, ASML, ASM International and Besi account for a total of approximately 23 % of the Netherlands' AEX index.

Figure 3: Weight of AI-linked firms in their national equity index (%)



Notes: Figures represent the combined index weights of the selected AI-linked companies. Reporting dates differ across indices, and the figures are therefore not strictly contemporaneous. Samsung Electronics includes preferred shares.
Source: Authors' elaboration based on data from KOFIA (KOSPI, September 2026 average), TAIFEX (TAIEX, 30 September 2026), SPDR S&P 500 ETF holdings (8 October 2026), Nikkei Inc. (Nikkei 225, 30 September 2026), and Euronext (AEX, 30 June 2026).

A correction in AI-related valuations could therefore have a simultaneous impact on several national equity markets. The combination of commercial interdependence and stock-market concentration makes them more sensitive to general changes in expectations regarding future AI demand and profitability.

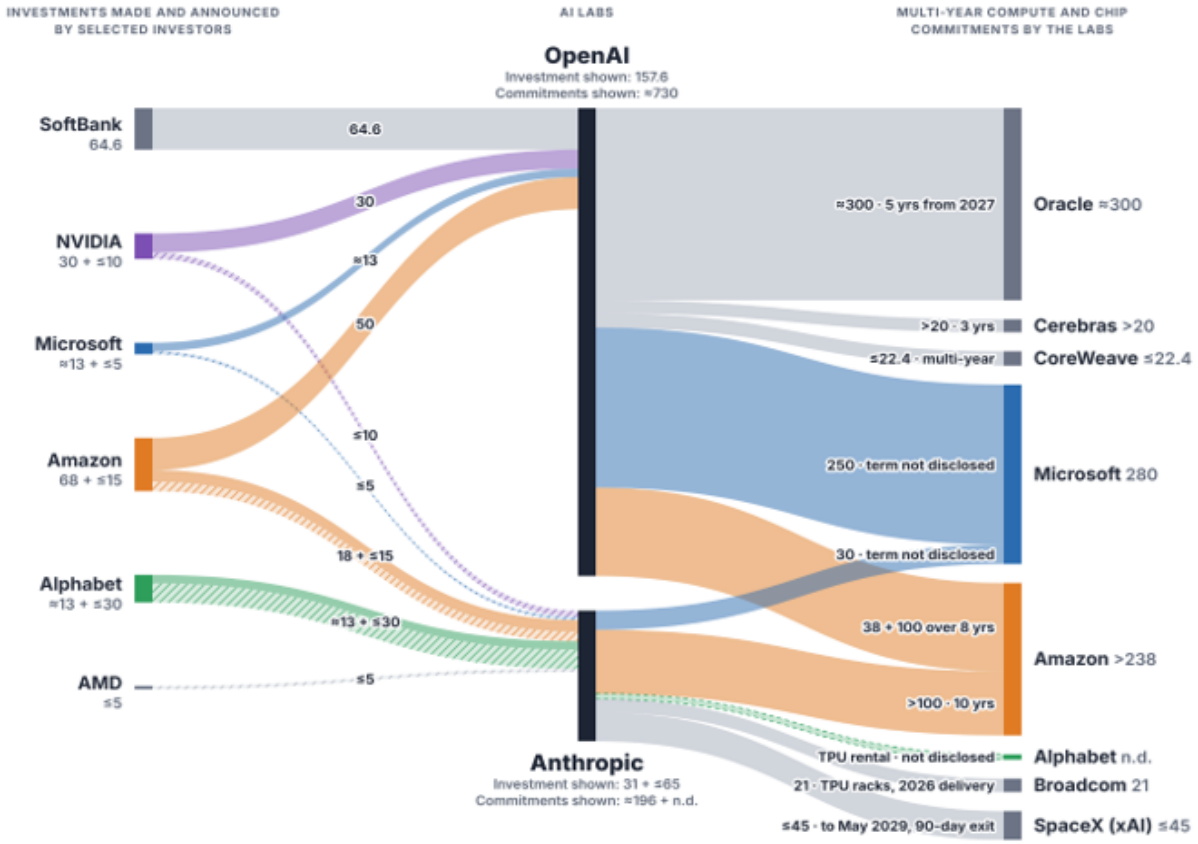
Concentration also has a competition-policy dimension. The financial concentration described above does not, by itself, establish an infringement of competition law. However, control of scarce inputs such as advanced semiconductors, cloud capacity, data and specialist expertise, combined with investments and commercial partnerships across the AI value chain, may create bottlenecks, increase switching costs or reinforce the market power of incumbent firms. These effects could strengthen the commercial and financial interdependencies described in this paper.

FINANCIAL CIRCULARITY STRENGTHENS COMMERCIAL DEPENDENCIES

Commercial interdependencies in the AI value chain are often supplemented by financial linkages between suppliers and customers. Key technology firms are not only supplying the infrastructure and computing power that AI developers require, but also funding them through equity investments, guarantees and other financial mechanisms. In return, AI developers enter into substantial commitments to purchase computing services and equipment from their financiers. This creates what can be described as financial circularity: capital provided by technology suppliers to AI developers supports investment and spending that, in turn, generates revenues for those suppliers. Such arrangements can help finance infrastructure development and secure demand for expensive computing assets, but they also reinforce the financial dependence of suppliers and customers on one another.

The scale of these relationships is considerable (see Figure 4). Amazon has invested around USD 50 billion in OpenAI and USD 18 billion in Anthropic (with up a further USD 15 billion conditional on compute-capacity milestones), while Microsoft has invested around USD 13 billion in OpenAI and Alphabet the same amount in Anthropic. NVIDIA has also committed substantial equity financing to OpenAI. Other arrangements involve warrants, financing vehicles and guarantees[1].

Figure 4: Investments in OpenAI and Anthropic and their multi-year compute and chip commitments (USD billion)



Notes: Solid flows represent investments made, while hatched flows indicate announced or conditional commitments. Matching colours identify the same company on both sides of the figure. Amounts include disbursements and future commitments with different contractual terms, time horizons and conditions; they are neither cash flows for the same period nor a comprehensive account of each company's financing. Symbols indicate approximate amounts (≈), amounts exceeding a stated value (>) and amounts up to a maximum (≤); n.d. indicates that the amount has not been disclosed. Selected investments and agreements announced up to 11 October 2026.
Source: Authors' elaboration based on company announcements and regulatory filings, and press reports.

On the purchasing side, OpenAI has committed approximately USD 250 billion to Microsoft's Azure cloud services and USD 138 billion to Amazon Web Services (AWS), while Anthropic has announced commitments exceeding USD 100 billion to AWS over ten years and USD 30 billion to Azure. These arrangements involve different contractual terms, time horizons and conditions, and should not be interpreted as directly comparable financial flows. Nevertheless,

[1] For example, AMD has granted warrants linked to chip purchases, while NVIDIA has provided guarantees associated with data-centre leases and committed to purchase certain unused computing capacity from CoreWeave.

they illustrate the scale of the financial and commercial relationships underpinning AI investment.

The implications of the equity stakes are far wider than the immediate interests of the equity buyers. The purchase commitments of OpenAI and Anthropic amount to substantial portions of the reported backlogs of major cloud providers, which underlines the degree of their commercial interdependence[2].

These arrangements can lead to the simultaneous occurrence of commercial and financial losses. If AI developers experience difficulties in the raising of capital or generate insufficient revenues, then their suppliers may be subject to both a reduction in demand for computing services and losses on their equity investments, financing arrangements or guarantees. In effect, financial circularity will amplify the transmission of risks across companies and make exposures harder to evaluate because commercial revenues and financial valuations are based on the same expectation about the future demand for AI.

THE GROWING ROLE OF DEBT, OFF-BALANCE-SHEET COMMITMENTS AND PRIVATE CREDIT

A third source of vulnerability arises from the way AI infrastructure investment is increasingly financed. The scale of investment in data centres, semiconductors and computing infrastructure has led to a rapid expansion of borrowing and long-term financial commitments. The IEA (2026) finds that the capital requirements of data-centre expansion have become too large to be financed solely through corporate balance sheets, increasing the importance of external funding and exposing investment plans to changing financial market conditions.

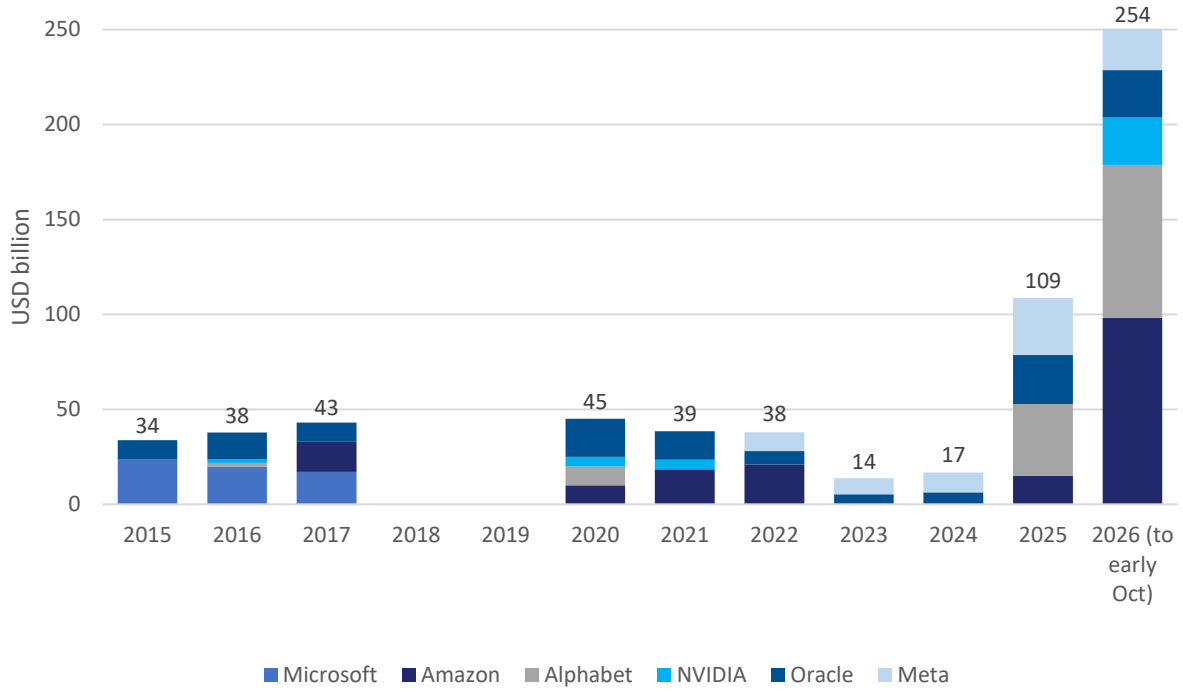
The scale of investment must also be assessed against the revenues required to provide an adequate return. Estimates are highly sensitive to assumptions about utilisation, pricing, financing costs, depreciation and the share of AI revenues accruing to infrastructure owners. In one scenario, data centre owners would need to generate approximately USD 8.5 trillion in cumulative revenue through 2030 to cover announced capital expenditure, financing costs and a 12 % return on invested capital, compared with no more than USD 3.6 trillion in attributable enterprise and consumer AI revenue under its forecast. These figures are scenario-dependent rather than a prediction of inevitable losses, but they illustrate the scale of the monetisation challenge. A recent Bank for International Settlements study similarly finds that competition for future AI revenues can produce over-investment when firms make investment decisions based on optimistic demand expectations.

Corporate bond issuance provides one indication of this shift (see Figure 5). Between 2015 and 2024, annual bond issuance by Amazon, Alphabet, Meta, Microsoft, Oracle and NVIDIA

[2] The announced purchasing commitments of OpenAI and Anthropic to AWS amount to almost USD 250 billion, equivalent to approximately half of AWS's reported backlog. OpenAI's reported USD 300 billion agreement with Oracle is equivalent to around 45 % of Oracle's backlog, while the two developers' commitments to Azure correspond to approximately 40 % of Microsoft's commercial backlog. These comparisons provide an indication of the scale of the relationships rather than precise measures of backlog concentration, since the commitments cover different periods and may not be fully reflected in reported backlogs.

averaged approximately USD 27 billion. It rose to USD 109 billion in 2025 and reached around USD 254 billion in 2026 through early October. The increase reflects both the exceptional scale of planned investment and companies' ability to access long-term financing at relatively favourable conditions.

Figure 5: Gross bond issuance by six AI-linked firms (USD billion)



Notes: Gross issuance of benchmark senior bonds by Amazon, Alphabet, Meta, Microsoft, NVIDIA and Oracle, by calendar year of pricing. Non-USD issuance is converted at the USD equivalent reported at pricing. Excludes commercial paper, term loans and borrowing through off-balance-sheet vehicles. No qualifying issuance was recorded in 2018 and 2019. Figures for 2026 cover issuance through early October.
Source: Authors' elaboration based on SEC term sheets and prospectus supplements, Cleary Gottlieb deal notes, International Financing Review (IFR) and Bloomberg.

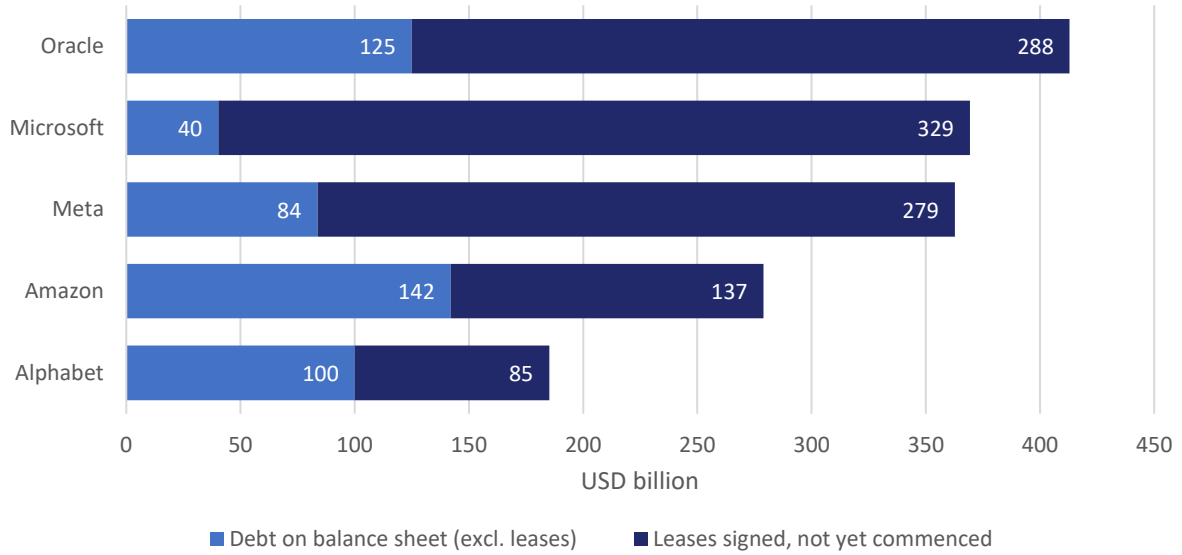
Investors have so far absorbed this issuance, including bonds with maturities extending to 50 years and, in one case, 100 years. Issuance in euros, sterling, yen and other currencies has also broadened the investor base beyond the United States. The growing supply of long-dated corporate debt could have wider implications for capital markets, including competition with sovereign issuers for long-term institutional investment.

However, traditional corporate bonds are only one part of the financing of the AI infrastructure. In addition to corporate bonds, companies increasingly use long-term leases, SPVs and private credit to finance data centres, and computing equipment. These structures divide financing responsibilities among different legal entities and investors, obscuring the overall picture more than just looking at the balance sheets of the giant technology companies would suggest.

Long-term leases allow technology companies to use data centres developed and financed by third parties. Such arrangements can involve significant future payment obligations, some of

which are not captured as lease liabilities. At Microsoft, Oracle and Meta, commitments under leases signed but not yet commenced substantially exceed reported financial debt. Microsoft reports approximately USD 329 billion in such obligations compared with USD 40 billion of debt, Oracle USD 288 billion against USD 125 billion and Meta USD 279 billion against USD 84 billion (see Figure 6)[3].

Figure 6: Debt on balance sheet and leases signed but not yet commenced (USD billion)



Notes: Debt figures exclude lease liabilities. Lease commitments represent undiscounted future payments under signed agreements that have not yet commenced and are therefore generally not recognised as lease liabilities. The two measures are not directly comparable and should not be interpreted as a combined measure of indebtedness. Amazon's commitments include all lease types.
Source: Authors' elaboration based on company 10-Q and 10-K filings (June-August 2026) and StockAnalysis.com for reported debt.

A broader Wall Street Journal analysis of nine technology companies estimated approximately USD 3 trillion in commitments not yet recognised on their balance sheets, including uncommenced leases and long-term purchase, supply and infrastructure commitments. These amounts are disclosed in financial-statement notes, but their different terms, time horizons and accounting treatment make it difficult to derive a consistent measure of aggregate leverage. They should not be added mechanically to reported debt because many are undiscounted commitments extending over several years.

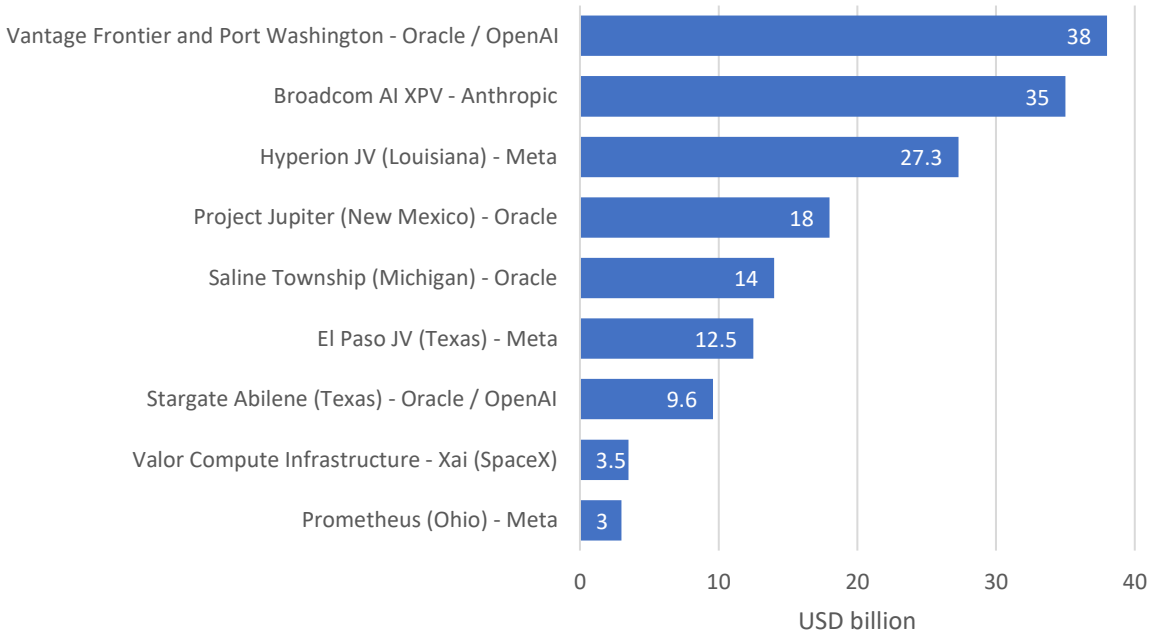
SPVs provide a further layer of financing. Instead of owning and financing the infrastructure themselves, technology companies can deal with separate entities that own data centres, raise financing and take payments from the technology company using the infrastructure. The borrowing arising from such arrangements may not consolidate on the balance sheet of the technology company, if the contractual arrangements and accounting treatment so provide,

[3] Under the applicable lease-accounting requirements, liabilities associated with signed leases that have not yet commenced are generally recognised when the underlying assets become available for use. The disclosed commitments represent undiscounted future payments over the contractual terms and are therefore not directly comparable with the outstanding book value of financial debt.

even though the technology company retains significant economic exposure via leases, guarantees or equity participation.

Meta's Hyperion data-centre project in Louisiana is an example. An SPV, owned 80 % by Blue Owl funds and 20 % by Meta, has issued USD 27.3 billion in bonds to finance the project (see Figure 7). Meta rents the campus, provides a residual-value guarantee of up to about USD 28 billion and reports a maximum exposure to the vehicle of roughly USD 46 billion[4]. Because Meta does not direct the activities that most affect the vehicle's performance, it does not consolidate it, and the debt does not appear on its balance sheet. Such transactions are growing significantly in scale: at least USD 161 billion in financings involving SPVs and project vehicles related to AI infrastructure have been identified.

Figure 7: Debt raised by SPVs and project vehicles linked to AI data centres and chips (USD billion)



Notes: Selected closed financing transactions involving SPVs and project vehicles linked to AI infrastructure, as of early October 2026. Amounts represent debt raised, rather than total project costs or equity investment. Transactions still under negotiation are excluded.
Source: Authors' elaboration based on company announcements and regulatory filings, Bloomberg, Financial Times and Data Center Dynamics.

Private credit is also playing a growing role, particularly in financing smaller cloud providers, data-centre developers and project vehicles. The share of AI investment financed by private credit rose from 9 % in 2024 to 34 % in 2025 (OECD, 2026). Outstanding private credit exposure to AI borrowers is estimated at over USD 200 billion, while projections suggest that private credit could provide substantially more financing for data-centre development over the coming years.

[4] Meta's reported maximum exposure includes its minority equity stake, lease commitments, future funding obligations and residual-value guarantee. These components represent different types of financial commitments and should not be interpreted as equivalent to outstanding financial debt.

These financing structures can facilitate investment through the mobilisation of capital outside the balance sheets of the largest technology companies. However, they also entail risks relating to leverage, refinancing requirements and valuations of the collateral, especially for smaller, more highly leveraged borrowers. For example, CoreWeave raised its outstanding debt from USD 21.4 billion as of the end of 2025 to USD 35.1 billion as of June 2026. Some of CoreWeave's debt is secured by graphics processing units (GPUs), the value of which can plummet quickly as new generations of chips come on-stream. If the financing tenors exceed the useful life or economic viability of the collateral, lenders can be exposed to greater risks of collateral and refinancing.

Early signs of pressure have also appeared in some financing vehicles. In September 2026, loans associated with an Oracle-linked data-centre project in New Mexico were reportedly quoted at around 90 cents on the dollar, and bonds financing Meta's Hyperion project traded at around 94 cents. These movements do not indicate a broad deterioration in AI credit markets, but they illustrate how shifts in expectations for the underlying project might reverberate into the prices being paid on such financing instruments.

The broadest concern is that AI-lending, commitments, and exposures are syndicated across corporates' balance sheets, forward lease commitments, project entities, private credit funds and other financial intermediaries. Some of these exposures are disclosed, but in different forms and with varying levels of granularity. It becomes challenging to evaluate the aggregated exposures of financial institutions to the same companies, projects or shared expectations about AI demand.

Taken together, concentration, financial circularity and increased reliance on debt create several channels through which the adjustment in AI investment boom can lead to financial loses. That adjustment's impact would depend on the size of the shock, ability of firms and investors to cope with the losses and the degree to which leveraging, liquidity effects, and common exposures amplify the adjustment. The above transmission channels will be discussed in the next section.

HOW AN AI CORRECTION COULD SPREAD THROUGH THE FINANCIAL SYSTEM

The financial risks posed by the AI investment boom could become systemic, even though the starting position is solid. The dominant technology firms remain highly profitable, have strong credit ratings and finance a large part of their investments with internally generated cash. Banks' direct exposure to these firms is also relatively limited. A correction in AI-related valuations would therefore not automatically lead to systemic financial stress.

The key question is how a large enough correction could spill over into the financial system. The concentration of AI-related investments, the growing use of debt and the financial linkages described in Section 2 suggest various potential transmission channels. A downward shock to expected AI revenues or valuations could result in losses for shareholders and creditors, reduced value of collateral for infrastructure financing and stress on leveraged borrowers.

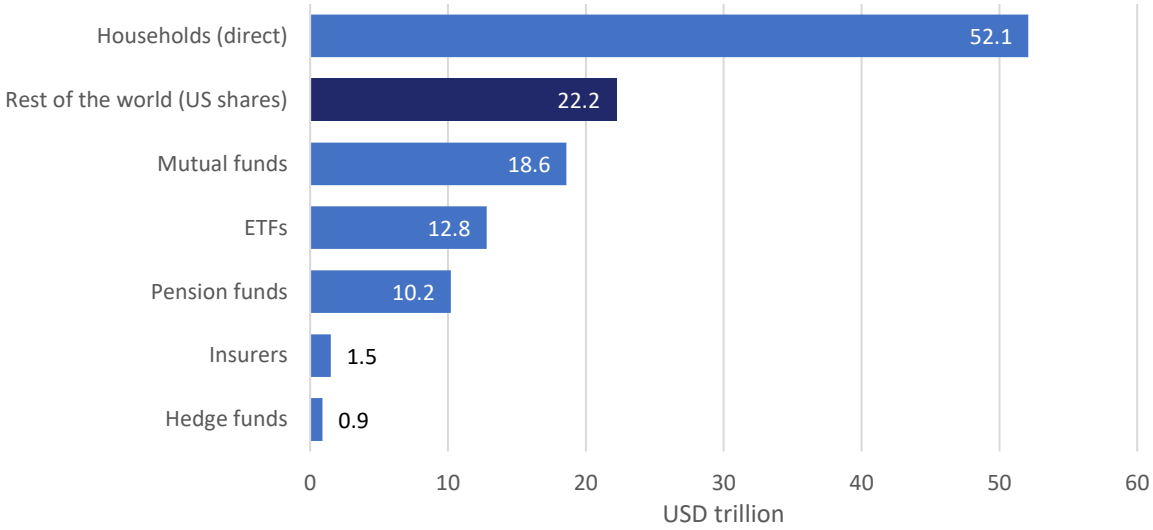
These losses would then reverberate through the financial system and markets, in particular if investors have similar positions, are affected by liquidity constraints or have recourse to short-term financing.

WHO BEARS THE FINANCIAL RISKS?

The risks of the AI investment boom go far beyond the tech sector itself. Shares and bonds of AI companies are owned by households, funds, insurers, pension funds and other institutional investors, both in the US and around the world. The financing of data centres and computer infrastructure is also provided by private equity funds, private credit and banks.

Equity investors are most exposed to a potential crash in AI values. US households own some USD 52 trillion in corporate shares directly, while mutual funds, exchange-traded funds (ETFs) and pension funds own a further USD 42 trillion in shares, mainly on behalf of households (see Figure 8). Large asset managers like Vanguard, BlackRock and State Street thus own significant stakes in the major AI companies, including NVIDIA, Microsoft and Broadcom.

Figure 8: Holders of corporate equities in the US financial accounts (USD trillion)

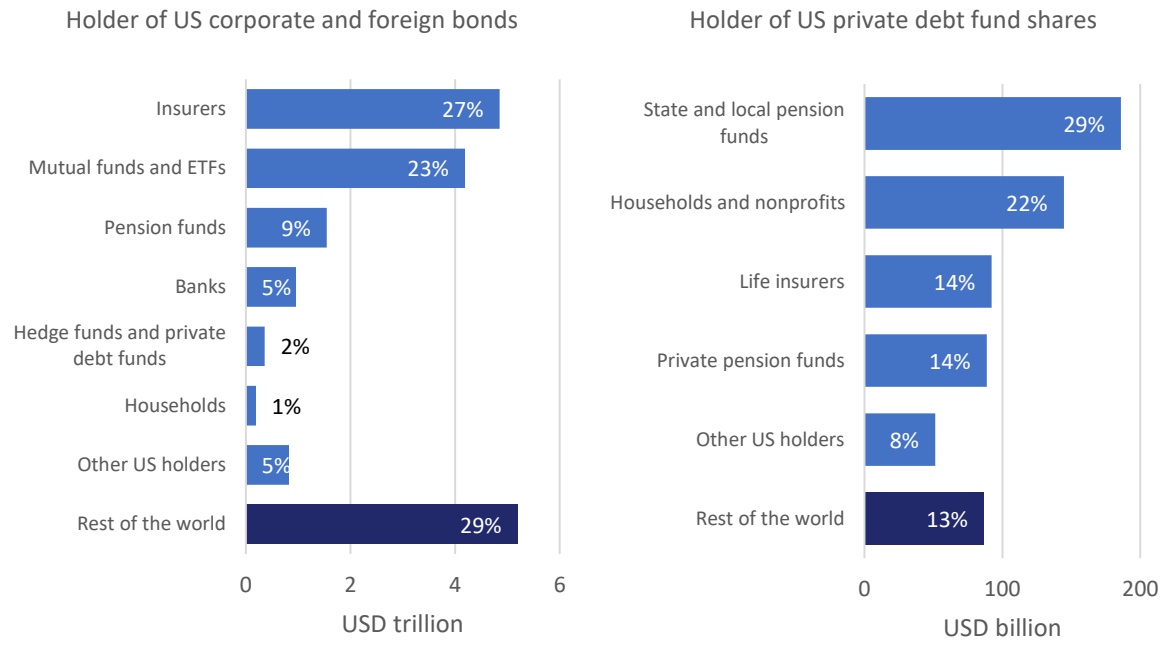


Notes: Figures cover corporate equity holdings across the US financial accounts, not AI-specific investments. Holdings of US residents include both US and foreign equities, while the 'rest of the world' category covers foreign investors' holdings of US equities only. Holdings through investment funds and pension funds may ultimately be attributable to households.
Source: Authors' elaboration based on Federal Reserve, Financial Accounts of the United States (Z.1), table F.51.1.s, Q2 2026.

These figures relate to the general ownership of corporate equities rather than specific AI investments. However, the figures help to understand how valuation losses in more concentrated equity markets can affect a broad investor base. Indeed, given the significant weight of AI-related companies in major equity indices, investors can be exposed to valuation losses related to the AI industry, even when investing in diversified funds and using passive investment strategies. The effects of such impact are not limited to the American market; as foreign investors own roughly one fifth of US equities, worth USD 22 trillion in total, a correction of AI-related valuations would therefore affect a global investor base.

Credit exposures are allocated to a different but partially overlapping group of investors. Of the approximately USD 18 trillion corporate bonds held in the US financial accounts, insurers account for 27 %, mutual funds and ETFs for 23 %, pension funds for 9 % and banks for 5 % (see Figure 9). Foreign investors hold approximately 29 %. The numbers relate to the wider area of the corporate bonds market rather than the AI-specific debt, but they illustrate the significant role of institutional investors in the absorption of corporate borrowing.

Figure 9: Who holds US corporate bonds and who invests in US private debt funds



Notes: The left panel depicts all corporate and foreign bonds held in the US accounts, not only those of AI firms (2026 Q2). The right panel depicts holders of shares in US private debt funds (2026 Q1).
Source: Authors' elaboration based on Federal Reserve, Financial Accounts of the United States (Z.1), table F.3.5.s (corporate and foreign bonds, Q2 2026) and table F.522.3.s (private debt funds, Q1 2026).

The growing issuance of AI-related corporate debt is increasing the relevance of these investor groups. According to the Bank of England, hyperscalers represented approximately 3 % of outstanding US investment-grade corporate debt at the end of 2025 but accounted for more than 15 % of issuance in the first part of 2026. Large institutional investors are therefore becoming increasingly exposed to the financing needs of the AI sector.

Beyond conventional corporate bonds, the financing of AI infrastructure through SPVs and private credit involves additional investors. Bonds issued by data-centre financing vehicles have been purchased by major asset managers and insurers. For example, PIMCO reportedly acquired approximately USD 18 billion of the USD 27.3 billion bonds issued for Meta's Hyperion project, while BlackRock-managed funds purchased more than USD 3 billion.

Private credit funds finance smaller activities in the cloud industry, data centre developers and other activities related to AI. The capital in private credit funds is mostly provided by institutional investors including pension funds and insurers, but also by households. In the US,

pension funds own about 42 % of private debt funds' shares, households 22 % and life insurers 14 %. In the euro area, insurers and pension funds contribute about 70 % of the capital in private credit funds (Cera et al., 2026).

Banks also participate in these financing structures, even if they do not directly lend to AI companies. They offer credit lines and other forms of financing to non-bank financial institutions, including private credit funds. The FSB estimates bank credit lines to private credit funds at roughly USD 220 billion for all sectors, while the lending by US banks to non-bank financial institutions amounted to USD 2.05 trillion in September 2026.

The same types of investors therefore emerge in different layers of AI financing: listed equities, corporate bonds, project-financing vehicles and private credit funds. The end investor exposure tends to be in the form of households' savings, investment funds, pension entitlements and insurance products. This does not imply that all these investors are necessarily exposed to the same assets or risk factors. What it does imply is that simultaneous losses on a number of different asset classes and financial institutions could occur as a result of a common deterioration in the prospects for AI investments.

HOW LOSSES COULD SPREAD: FROM VALUATIONS TO CREDIT AND LIQUIDITY PRESSURES

A correction in the expectations for investments in AI could first manifest itself in falling share prices. Reduced expectations of revenues or profits would lower the valuations of not only the AI developers but also the technology suppliers and infrastructure enablers. As a result of their strong weighting in major equity indices, losses could ripple through stock markets and portfolios, depressing wealth of households and institutional investors. In addition, firms that hold stakes in AI developers could see their investments shrink in value, and financial circularity could amplify the losses if suppliers simultaneously experience lower valuations and reduced demand for their products and services.

The financial spillovers could spread beyond equity markets if reduced demand for AI computing capacity impacted on borrowers' and financing vehicles' ability to repay their obligations. **The first channel of credit transmission relates to revenues and refinancing capacity of the AI-related borrowers**. Borrowers that have borrowed money to finance data centres and computing equipment rely on future revenues from AI services or long-term contractual payments. Falling demand short of expectations may result in borrowers' inability to meet obligations such as servicing their debt, and/or refinancing outstanding obligations or new financing.

Such risks will be of particular relevance to smaller and more leveraged companies. Unlike the biggest names in the technology sector, who tend to have ample cash flow and strong balance sheets, some infrastructure providers are dependent on external financing. A reduction in their projected revenues or access to the capital markets, therefore, could quickly lead to financial stress.

A second channel relates to the value of the assets which secure borrowing relating to AI. Data centres and graphics processing units (GPUs) are increasingly used as collateral in financing arrangements. The value of these assets is partly determined by expectations about demand for computing capacity and, in the case of GPUs, improvements in technology that entail a shorter economically useful life.

A decline in demand for AI could therefore affect the revenues from such assets, as well as their collateral value. In situations where loans are taken out against assets that lose value more quickly than the debt owed against them, creditors are likely to experience significantly greater losses if there is a default on the loan. The BIS has also warned that the same underlying assets may be pledged to several financing vehicles, which could complicate the identification and distribution of losses.

A third channel derives from the contractual nature of data-centre financing instruments. Certain SPVs finance the infrastructure by issuing long-dated debt while relying on a lease agreement with technology companies that might have shorter terms of initial lease or renewal. Failure to exercise a right to renew may lead to the inability of the SPVs to cover their debt servicing costs given the lower demand for computing power.

In such cases, creditors may rely on the re-sale value of the infrastructure or on guarantees of the residual value provided by the technology companies. These guarantees could be called upon should the values of the assets decline, and this would transfer the financial stress back to the companies that underpin the projects. Accordingly, the question is whether such losses will actually be incurred, depending on the precise nature of the contractual terms, and the financial ability of the parties.

A fourth channel relates to the financial institutions and markets. Losses in the value of equity, bonds or private credit investments will have an impact on investment funds, insurers and pension funds. Where investors suffer redemption requests, liquidity needs, or face demands for financing, they may be forced to sell assets or shrink in other investments (Coval and Stafford, 2007).

Such adjustments could have implications for other sectors. Redemption pressures in investment funds could result in the sale of liquid assets unrelated to the AI, leveraged investors could unwind their positions in order to meet margin requirements or financing commitments (Brunnermeier and Pedersen, 2009) and banks with lending commitments to affected non-bank financial institutions could experience losses or reduce their lending commitments. These adjustments could contribute to the initial shock by exacerbating downward pressures on asset prices and tightening financing conditions (Greenwood et al., 2015).

The extent to which these effects would take place would depend on the concentration of exposures, the leverage and liquidity positions of the affected institutions, and the ability of markets to take in asset sales. The more important risks would result if credit losses and

liquidity pressures lead to financial institutions transmitting the shock to other markets and economic activities.

WHEN COULD FINANCIAL VULNERABILITIES BECOME SYSTEMIC?

Whether the above transmission channels give rise to system-wide financial instability depends on the extent to which financial institutions and investors have concentrated exposures to similar sets of AI-related assets and counterparties, the refinancing and leverage requirements of the affected borrowers and intermediaries, and the ability of financial institutions and markets to withstand losses without having to sell off assets or contract their lending activities. These factors will determine whether losses are confined to the set of investors and institutions with concentrated exposures or will spill over into other financial markets and economic activity.

The evidence presented does not allow a systemic risk to be ruled out: exposures are large, growing and interconnected, and the data needed to measure them are incomplete. To assess the potential for systemic implications, more information is needed on the distribution of exposures and the ability of affected institutions to absorb losses and on the interconnections between financial intermediaries. The implications for European financial markets and supervisory frameworks are discussed in the next section.

WHY EUROPE IS EXPOSED AND WHAT IS ALREADY BEING DONE

The global financing of the AI investment boom creates several potential channels of exposure for European investors and financial institutions. This section aims to explore those exposures, consider how European authorities are already tracking related weaknesses and see whether the current information is enough to assess their financial stability implications.

EUROPE'S EXPOSURE TO THE AI INVESTMENT BOOM

The most obvious transmission channel would be **Europe's significant exposure to US stock markets** in general and the big tech stocks in particular. According to the ECB, euro area investors own some USD 6 trillion of US equities, about four times more than ten years ago. Much of this is invested in funds, which are domiciled in Ireland and Luxembourg. As a result, a significant part of the European investments in US equities may take the form of exposure to AI valuations through diversified funds and index-tracking products, although it is difficult to say how big this exposure is based on the overall size of the US equity holdings of euro area residents.

The ECB has pointed out this fragility. In the May 2026 Financial Stability Review, it warns that euro area investment funds are the largest domestic investors in US equities and that these portfolios are exposed to assets with inflated valuations. A correction in AI-related share prices could therefore spill over into European investment funds and their investors, with possible repercussions in the broader market if redemptions or portfolio shifts were to take place on a large scale.

A second channel is **corporate debt**. European investors hold substantial amounts of US corporate bonds, including debt issued by major AI-related companies. Foreign investors held approximately USD 4.8 trillion of US corporate bonds in June 2025, with around 40 % attributed to euro area countries, particularly Luxembourg, Belgium and Ireland. EU insurers and pension funds held approximately EUR 262 billion of US corporate bonds.

The growing issuance of debt by major technology companies is also directly affecting European corporate bond markets. US hyperscalers increasingly issue bonds denominated in euros and other European currencies, attracting investors seeking high-quality, long-maturity assets. According to the ECB, these companies accounted for approximately 15 % of the increase in euro area holdings of euro-denominated corporate bonds in the year to March 2026, with particularly strong demand from pension funds and insurers (Duquerroy et al., 2026).

This development has two implications. First, European institutional investors are beginning to see more direct exposure to the financing of AI-related infrastructure investment. Second, the increased supply of long-dated corporate bonds could have more general implications for the financing environment. Large volumes of issuance compete for the same balance-sheet capacity of institutional investors, with possible repercussions for the yields and spreads of other corporate and sovereign debt. The ECB finds no such spillovers in the euro area so far, but this reverse crowding out is a forward-looking risk: it could intensify as hyperscaler borrowing rises, at a time of high public debt, rising long-term yields and less central bank buying.

A third channel is **private credit** and infrastructure financing. European insurers, pension funds and investment funds participate in these markets, including via cross-border investment structures. Their exposure to AI-related borrowers and financing vehicles could lead to losses if the investments' performance is poor, although the scale of such exposures cannot be established from the available data.

The above figures must be carefully interpreted: aggregate European holdings of US equities, corporate bonds or private credit are not proxies for AI-specific exposure. Fund domicile does not necessarily reflect the nationality of ultimate investors, and the entities that bear the economic risk (Beck et al., 2024). The evidence highlights the relevance of channels through which an AI correction might impact Europe, but not the overall size of European exposure to the AI investment boom.

A fourth channel is the **industrial value chain**. European corporations, most prominently semiconductor equipment manufacturers like ASML, have some exposure to investments in advanced computing and semiconductor manufacturing. A worldwide slowdown in AI investment would therefore affect European corporate revenues and equity values, independently of a direct financial link to US technology firms. The primary financial stability risk, however, is that the various channels might cause correlated losses in European markets and systems.

WHAT EUROPEAN AUTHORITIES ARE ALREADY DOING

European authorities are not starting from scratch. Financial stability monitoring and prudential supervision already cover several of the vulnerabilities associated with the AI investment boom, particularly market concentration, credit risk, leverage, liquidity mismatches and interconnectedness between banks and non-bank financial institutions.

The ECB has increasingly examined the financial stability implications of AI-related market developments. Its November 2025 Financial Stability Review highlighted the growing concentration of US equity markets in a small number of large technology companies and the potential for spillovers across markets. In May 2026, the ECB also drew attention to the increasing reliance of AI-related companies and infrastructure on credit financing, warning that the combination of rapid equity valuation growth and rising business debt warrants monitoring.

The ECB has also examined the implications of hyperscaler bond issuance for euro area financial markets, including growing institutional investor exposure, competition for long-term financing and the effects of bond index concentration. These assessments provide an important starting point for understanding how developments in the AI sector may affect European asset valuations and financing conditions.

The European Securities and Markets Authority (ESMA) monitors risks in securities markets and investment funds, including equity market concentration, liquidity, leverage and private finance. Its report dated March 2026 on trends, risks and vulnerabilities shows the dependence of equity funds' performance on the state of the US market and the risks associated with the development of private finance, primarily its opacity and interconnectedness (ESMA, 2026).

The European Systemic Risk Board (ESRB) addresses risks that non-bank financial intermediation poses to financial stability, and its linkages to the banking system. Its September 2025 risk monitor highlights leverage, liquidity mismatches, interconnectedness and data gaps as important vulnerabilities (ESRB, 2025). The analysis applies to private credit funds and financing vehicles which participate in AI infrastructure investment, though the ESRB's analysis is not specific to AI.

The European Insurance and Occupational Pensions Authority (EIOPA) also monitors investment risks for European insurers and occupational pension funds. In its December 2025 Financial Stability Report it investigated the growth of private credit exposures for insurers and institutions for occupational retirement provision. A subsequent factsheet published in July 2026 reported that private credit represented around 5 % of EEA (re)insurers' total assets at the end of 2025, with private equity representing a further 6.3 % (EIOPA, 2026). These figures relate to private asset investments across sectors not specific to AI financing, but they show that these areas are already covered by supervisory monitoring, representing a proportion of the potential investor base exposed to AI-related private credit.

Existing regulatory and reporting requirements also provide additional information such as prudential data relating to credit exposures, concentration and counterparty risks reported by bank supervisors, investment holdings and associated risks reported by insurers under Solvency II, and regulatory reporting and risk-management requirements that apply to investment funds, including by virtue of the Alternative Investment Fund Managers Directive (AIFMD). Financial reporting requirements also reveal information about recognised liabilities, leases and certain commitments and guarantees.

The Market Integration and Supervision Package (MISP) is also relevant to the coordination of asset-management supervision. The Commission proposal would require ESMA to conduct at least annual reviews of large cross-border groups of UCITS management companies and AIFMs, consolidating relevant information already held by national authorities and assessing supervisory approaches to group governance, resource allocation and risk management. The Council's political compromise instead relies on coordination colleges involving ESMA and the relevant national authorities. These arrangements could improve the group-wide assessment of large cross-border managers, but neither model would, by itself, provide a complete overview of AI-related portfolio exposures across the asset-management industry. That would still require sufficiently granular fund-level information and the ability to link common underlying borrowers, projects and financing vehicles across datasets.

These arrangements already provide useful information to European authorities on matters concerning the assessment of AI-related financial risks. However, they were not necessarily designed to capture a full range of commercial and financial relationships as described in this paper. They are contingent on the granularity of the information that is reported, identification of underlying counterparties and assets and the ability of different authorities to combine the information into a single picture across institutions and markets.

An important distinction is that the current EU financial regulation and supervision covers primarily the financial institutions and activities involved in financing AI investment, rather than the AI companies and infrastructural projects themselves. Banks, insurers and investment funds are subject to their own prudential, reporting and risk-management requirements, but many of the underlying borrowers and financing vehicles are based outside the EU. Consequently, the challenge is not only to monitor risks within regulated institutions, but also to understand their exposures to companies, projects and contractual arrangements that may be beyond the direct supervisory reach of European financial authorities.

WHAT US AUTHORITIES ARE DOING

US authorities have also begun to examine elements of the emerging risk. The Federal Reserve's May 2026 Financial Stability Report identified AI and private credit among the salient risks raised by market participants and examined the growth of bank commitments to non-bank financial institutions. Subsequent Federal Reserve work has examined these issues through separate strands of analysis. A May 2026 speech by Governor Lisa Cook considered the increasing use of corporate bonds, private debt and asset-backed credit to finance AI

infrastructure. A July 2026 Federal Reserve note developed indicators for monitoring the AI investment buildout, while later analysis examined the structure and interconnectedness of private-credit and leveraged-loan markets. The Federal Trade Commission has separately examined investments and partnerships between major cloud providers and AI developers, focusing on access to key inputs, switching costs and the exchange of competitively sensitive information.

These initiatives address important parts of the picture, but they remain distributed across financial-stability monitoring, prudential supervision, securities disclosure and competition policy. There is not yet a publicly available consolidated assessment linking exposures across listed equities, corporate bonds, uncommenced leases, purchase commitments, SPVs and private credit. This fragmentation reinforces the case for cross-agency analysis in the US and international cooperation through the FSB.

WHERE THE MONITORING GAPS REMAIN

The main issue is not necessarily the absence of regulation or supervisory reporting, but rather the dispersion of information across companies, financial instruments, institutions and jurisdictions. Exposures to the same AI-related borrowers or projects can be reported separately in company financial statements, investment fund disclosure and prudential reporting. The understanding of the potential transfer of losses through these exposures requires linking this information across different instruments and financial institutions (ESRB, 2025).

Three issues are particularly relevant. **First, identifying aggregate exposures to common counterparties and projects**. Supervisors may have knowledge about individual securities, borrowers or funds of investment without being able to trace the entire financial relationships between them. For example, direct holdings of bonds issued by a technology company can be identifiable, while indirect exposures (through private credit funds, securitisations or project-financing vehicles) could need more information.

This is important since several investors may have a common ultimate exposure to the same project or borrower through the use of different investment vehicles, and an insufficient knowledge of these relationships can result in an under-estimation or difficulty in measuring the level of aggregate concentration.

Second, assessment of contingent obligations and risks alongside those which relate directly to financing through debt. Long-term leases, residual value guarantees, purchase commitments and other contingent arrangements may generate significant off-balance-sheet exposures representing economic obligations poorly captured by conventional debt measures. While the publication of relevant obligations and commitments is part of many accounting frameworks, the disclosure criteria are subject to interpretation, and different presentation options and levels of detail may be adopted.

The issue is therefore not simply one of disclosure of these obligations but rather whether there is enough information to enable investors and supervisors to make consistent comparison and assessment of their implications for interlinked firms and financing vehicles. The challenge for European authorities is greater where the underlying companies are located outside the EU and subject to different reporting and supervisory regimes.

Third, the relationships between banks and non-bank financial institutions. Private credit funds, insurers, pension funds and investment vehicles can have exposure to AI-related borrowers, while banks can be the source of finance for some of these intermediaries. A shock to the underlying investments could then propagate losses across many institutions.

European authorities already supervise leverage, liquidity and interconnectedness in non-bank financial intermediation. However, determining the degree to which exposures across different types of institutions ultimately depend on the same AI-related investments may require aggregating supervisory information that is gathered under different frameworks.

The cross-border aspect creates additional complexities. The head offices of many of the largest AI companies, financing vehicles and underlying investments are based outside the EU where their securities and structures are held by investors in several jurisdictions. As a consequence, European supervisors may have knowledge about the exposures of the institutions subject to their authority, but not necessarily on the corresponding information on the borrowers, contract commitments or financing structures.

POLICY PRIORITIES: STRENGTHENING TRANSPARENCY, MONITORING AND FINANCIAL RESILIENCE

The analysis highlights three policy priorities: enhancing international transparency and collaboration, strengthening the identification of interlinked AI-related financial exposures, and introducing these vulnerabilities to financial stability assessments. The international layer takes precedence, as several of the borrowers, projects and financing vehicles are headquartered outside the EU, and no individual jurisdiction has all the pieces of data. The response should build on existing supervisory and regulatory pillars, and additional requirements could be considered if particular gaps are identified.

International coordination is needed not only because the relevant data are dispersed across jurisdictions, but also because the losses from an AI-related correction would not be confined to the US. Foreign investment funds, insurers, pension funds and households hold US equities, corporate debt and private-market exposures, meaning that a correction centred on US companies and infrastructure could transmit losses globally.

STRENGTHEN INTERNATIONAL TRANSPARENCY AND SUPERVISORY COOPERATION

The G20 should encourage the FSB to take up the topic in its work on private credit and non-bank financial intermediation, given the cross-border nature of AI financing. There are three things it could consider.

First, improve comparability of financial disclosures. The International Accounting Standards Board (IASB), the Financial Accounting Standards Board (FASB) and relevant securities regulators should assess whether to strengthen existing requirements to ensure sufficiently comparable information about material lease commitments, residual-value guarantees, capacity backstops, warrants and long-term purchase commitments. If so, targeted improvements could be made, with the support of the International Organization of Securities Commissions (IOSCO), to promote international co-ordination. In the EU, ESMA and national accounting enforcement authorities could review the consistency and completeness of such disclosures.

Second, enhance cross-border supervisory data sharing. The FSB in cooperation with relevant international standard-setting bodies could encourage greater consistency in reporting and promote the exchange of information necessary to identify exposures to common AI-related borrowers, projects and financing vehicles, subject to applicable legal and confidentiality requirements. Particular attention should be paid to exposures held through private credit funds, securitisations and SPVs.

Third, promote internationally coordinated scenario analysis. The FSB and IMF could devise common assumptions and methodologies for assessing AI-related financial shocks and cross-border transmission channels that could inform national and regional financial stability assessments, without requiring uniformity in supervisory exercises across jurisdictions.

IMPROVE THE IDENTIFICATION AND MONITORING OF AI-RELATED FINANCIAL EXPOSURES

European authorities should evaluate whether current supervisory data captures direct and indirect exposures to AI-related companies and financing structures. The ECB, the European Supervisory Authorities (ESAs) and the ESRB, in cooperation with national competent authorities, should look at exposures to alternative forms of investment beyond traditional equity and bond holdings, such as SPVs, securitisations, private credit funds and other AI infrastructure financing vehicles.

The priority should be to identify common underlying borrowers, projects and collateral across different financial instruments and institutions, enabling supervisors to detect concentrations which might not be apparent from individual portfolio analyses. There is a need to assess whether existing sets of supervisory data across different sectors could be linked together, and whether the detail is available to enable identification of underlying exposures. The exercise should start with an assessment of existing reporting frameworks, identifying what data are available and gaps in their coverage and accessibility, before identifying any additional requirements.

Responsibilities should be based on existing mandates. The ECB and national banking supervisors would analyse banks' direct and indirect exposures, including banks' links to non-bank financial institutions, and EIOPA and national insurance and pension supervisors would assess institutional investors' exposures, as would ESMA and national securities regulators for

investment funds and securities markets. The ESRB could provide an overarching assessment of concentration and interconnectedness based on these findings.

The MISP trilogue provides an immediate opportunity to strengthen this coordination. The final framework should retain a meaningful role for ESMA in the periodic assessment of large cross-border asset-management groups and ensure that ESMA can obtain and consolidate comparable information from the relevant national authorities. Coordination colleges can support this process, but they should not become a substitute for an independent group-wide assessment. These reforms would strengthen the institutional framework for identifying concentrated and interconnected exposures, although additional work would still be needed to identify common underlying assets across individual funds and financing structures.

A related issue concerns competition policy. Competition authorities should continue to scrutinise whether partnerships, exclusivity arrangements, vertical integration or control of scarce inputs restrict market access, reinforce lock-in or allow market power to be extended across the AI value chain. This role is distinct from financial-stability supervision: competition authorities assess the structure and functioning of markets, while financial supervisors examine concentrations of exposure and channels of loss transmission. Nevertheless, the two perspectives should inform one another where the same corporate relationships create both market power and financial interdependence.

INTEGRATE AI-RELATED VULNERABILITIES INTO FINANCIAL STABILITY ASSESSMENTS

The European authorities should consider including AI-related vulnerabilities in their assessments of financial stability and, where relevant, stress-testing exercises. A common analytical scenario could involve assuming a correction in AI-related equity valuations together with decreased demand for computing infrastructure, lower collateral values and refinancing pressures on leveraged borrowers and project vehicles. The scenarios could also reflect the potential propagation of losses through private credit funds, insurers, pension funds, investment funds and banks.

The objective would be to identify vulnerabilities across institutions and asset classes, rather than to predict an AI-driven financial crisis. The ECB, ESRB and ESAs could collaborate on a common analytical framework, which individual supervisors would then adopt in line with their mandates. Special focus should be placed on bank to non-bank interconnections, liquidity pressures and the risk of forced asset sales not only in the AI sector.

For European policymakers, the key issue is to ensure that existing regulatory and supervisory frameworks are fit to capture and assess interconnected financial exposures from AI investment (including risks from outside the EU) and not that a separate regulatory regime is needed. Improvement in the use of data by supervisors, more comparable disclosures and enhanced international cooperation would support financial resilience without unduly constraining investment in AI.



www.ecmi.eu



ECMI - European Capital Markets Institute



@ECMI_CEPS

