Initial conditions determine upside risks to inflation[1]

In basis points

Graph 9

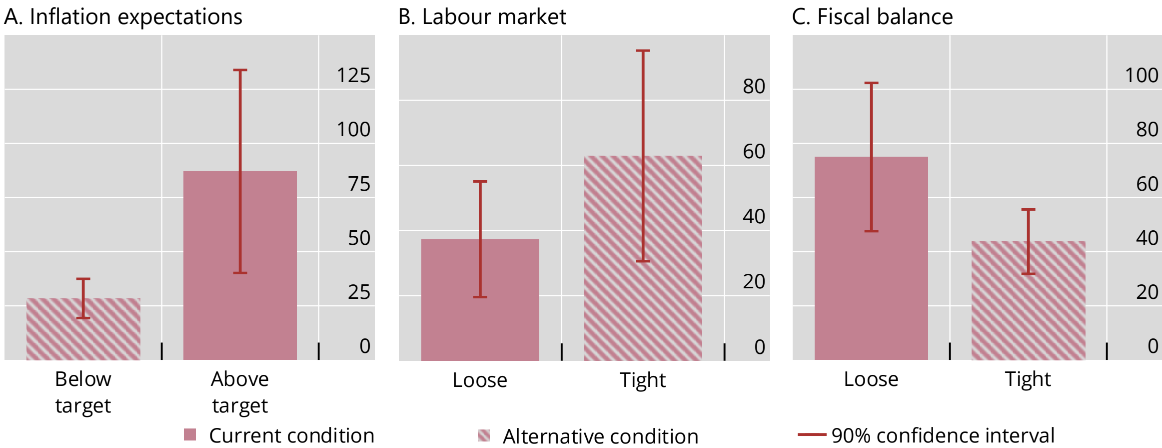

[1] Impact of a 10% oil supply shock on headline inflation after 12 months for respective categories of initial conditions. Oil price growth instrumented with oil supply shocks identified following Baumeister and Hamilton (2019). See additional notes to graphs for details.

Sources: Baumeister and Hamilton (2019); OECD; IMF; Consensus Economics; LSEG Datastream; Macrobond; national data; BIS.

Several mitigating factors could help contain the inflationary impact of the current oil shock. First, the prospective resumption of oil flows through the Strait, if sustained, should help truncate the upside tail risk to inflation. Another is the reduction in energy intensity, which has fallen by more than half in some countries over the past two decades. This reduces the sensitivity of domestic economic activity to higher oil prices, although these gains may be tempered by the embedded energy costs of imported goods. Another mitigating factor is the more favourable macroeconomic context relative to previous inflation surges. Inflation expectations, while still above target in some jurisdictions (Graph 9.A), are lower now than after the start of the Russia-Ukraine war in 2022. Labour market normalisation since 2022 has further reduced the risk of second-round effects and wage-price spirals (Graph 9.B). And while fiscal balances remain looser than average (Graph 9.C), relief spending that is more targeted this time should help strengthen fiscal discipline and lessen the inflationary impact of the current energy shock.

AI progress and investment boom under pressure

AI has the potential to raise productivity significantly over the coming decade. Task-level studies consistently report large efficiency gains, often to the tune of between 20 and 50% in time savings (Graph 10.A). Aggregate productivity growth estimates tend to be more conservative at less than 1% over a long horizon, reflecting challenges in adopting the technology at scale and integrating it with production processes. Still, there are further upside productivity gains, particularly if the technology improves to the point at which knowledge creation can be automated. The potential implications of such transformative AI for growth, income distribution and monetary policy are profound (Box C).

Box C

Transformative AI, long-term growth and r-star

Artificial intelligence (AI) has the potential to differ fundamentally from earlier waves of technological progress. Previous general purpose technologies, such as the steam engine, electrification and information technology, raised workers' productivity by providing them with better tools. AI could go further by augmenting the production of knowledge itself. If, at some point, AI systems can improve their own capabilities and "create" technology and ideas, the macroeconomic consequences could be profoundly different from past innovations. A key constraint on long-run growth, namely the rate at which humans can generate new ideas, could be lifted.

Recent research has begun to formalise this possibility. Trammell and Korinek (2023) argue that if AI capital becomes a sufficiently close substitute for human labour, the economy can transition from the regime of constant exponential growth to one of accelerating super-exponential growth. As machines autonomously improve themselves, the economy acquires a self-reinforcing engine of growth. Jones and Tonetti (2026) identify a supply side counterforce. If different economic tasks are strong complements, the overall output would be constrained by the task that is improving slowest, the weakest link. AI progress in automated tasks could then fail to translate into faster aggregate growth if essential human performed tasks hold output back. Trammell and Korinek (2023) also argue that economies currently operate in this complementarity regime, implying that the transformative scenario remains a possibility rather than a present reality.

These frameworks are supply side oriented – they focus on the productive potential of AI and assume that consumer demand keeps pace with supply. But as AI advances, automation increasingly diverts income from labour, which is spent on goods and services, into further AI investment. The consumer base could erode as productive capacity expands. Forward-looking firms, recognising the shrinking future market for AI-produced goods and services, may find it unprofitable to invest in innovating and automating the next task. Productivity stalls not because of technological limitation, but because the demand to justify further capacity expansion is missing. The demand bottleneck becomes the binding constraint.

AI's long-run impacts are uncertain and scenario-dependent[1]

Graph C1

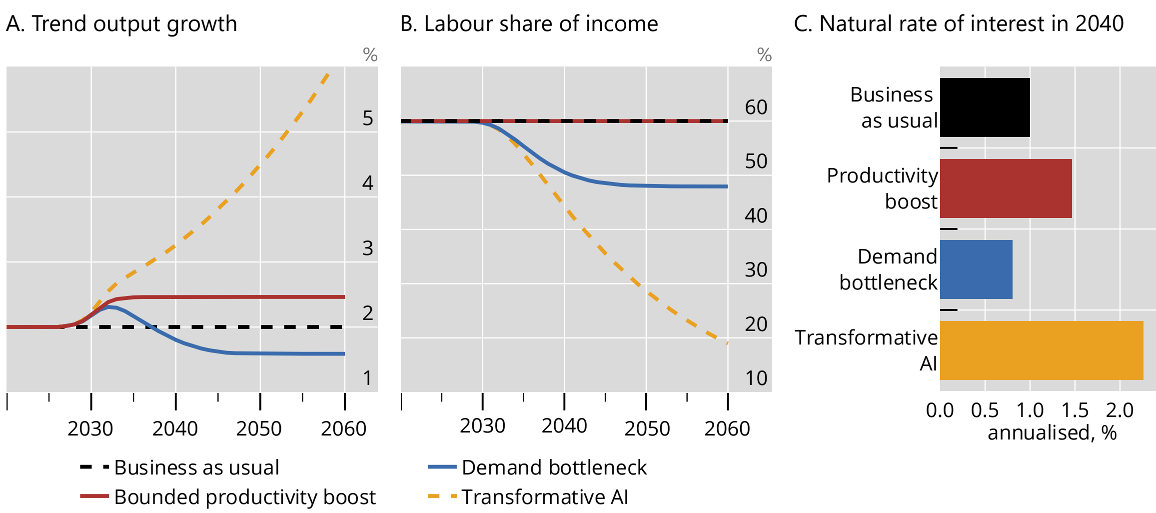

AI = artificial intelligence.

[1] Simulation based on a task-based growth model featuring potentially transformative AI capital and two types of households with different marginal propensities to consume.

Sources: Rungcharoenkitkul (2026a); BIS.

A unified framework that nests the supply side predictions as special cases, but incorporates demand bottleneck, synthesises the range of possibilities and provides insights for monetary policy. The framework features a task-based production structure in which AI capital can automate a rising share of tasks. It has two classes of households (capital owners and workers) to track who bears the costs of adjustment[1] and features a forward-looking free-entry condition governing firms' decisions to invest in further innovation and automation. In calibrating the model, trend productivity, markups, product turnover and longevity of innovation rents are set to standard empirical benchmarks. The model economy's labour share and r-star before the introduction of transformative AI are matched to data and established estimates. By varying the scope of automation and innovation's reliance on consumer demand, the model generates scenarios that include: (i) business as usual; (ii) bounded productivity gains; and (iii) explosive growth. Critically, it identifies a fourth possibility, the demand bottleneck.

The projections for aggregate output highlight the wide uncertainty surrounding AI's growth impact (Graph C1). Under the business as usual scenario in which labour remains the binding input, output continues to grow at its historical trend rate of 2% annually. A bounded productivity boost would shift trend growth permanently by a constant margin – which has not occurred since the Industrial Revolution. Under the transformative AI scenario, in which AI becomes a self-reinforcing growth engine, output growth expands exponentially and the labour share falls towards zero. Under the demand bottleneck scenario, output growth rises before falling below its historical trend in the long run as automation stalls. Since every displaced worker is also a lost consumer, the spending that rewards innovation eventually shrinks, leading to a slowdown. This scenario produces a labour share decline that is less severe than the transformative case, but worse than the other two cases.

The analysis also sheds light on the implications for the natural rate of interest, or r-star. Under the business as usual scenario, r-star stays at its pre-AI level. Under the bounded productivity boost and transformative AI scenarios, higher productivity growth raises the marginal product of capital and r-star increases. Under the demand bottleneck scenario, r-star initially rises while supply side forces dominate but then falls below its pre-AI baseline as the bottleneck takes hold. The r-star implications thus depend on parameters that are uncertain today: how much AI profits and innovation rely on consumer demand, how quickly competition within the AI sector erodes margins, and how fast AI products and infrastructure become obsolete. The corresponding medium-term inflation pressures move with r-star in direction, turning disinflationary if the demand bottleneck dominates and inflationary if expected productivity lifts demand on net.

[1] See also Fornaro and Wolf (2026), who show that automation-driven redistribution towards capital owners with lower marginal propensities to consume can open a demand shortfall and a liquidity trap in a New Keynesian model, and Falk and Tsoukalas (2026), who show that a similar demand externality can lead firms to over-automate relative to cooperative equilibrium.

The transition to a more productive AI-driven economy entails risks, however. As more capable AI tools find applications in more tasks and occupations, labour displacement could intensify. Whether or not AI advances create new jobs – or expand demand for existing ones – sufficiently to make up for such displacements remains uncertain. Unlike past general purpose technologies, AI competes directly with human cognitive abilities, possibly narrowing the scope for workers to move up the value chain or find new non-disrupted tasks. To date, such disruptive labour displacements have yet to occur at scale. But there are signs of possible adjustments to come. In earnings calls, more firms are acknowledging potential productivity gains from AI, signalling their intent to automate an increasing share of production processes and engage in labour substitution (Graph 10.B). Consistent with this, US sectors with higher exposure to AI have also seen higher productivity gains, partly at the expense of lower employment growth relative to other sectors (Graph 10.C).

AI holds promise of large productivity gains, but with risks to labour markets[1] Graph 10

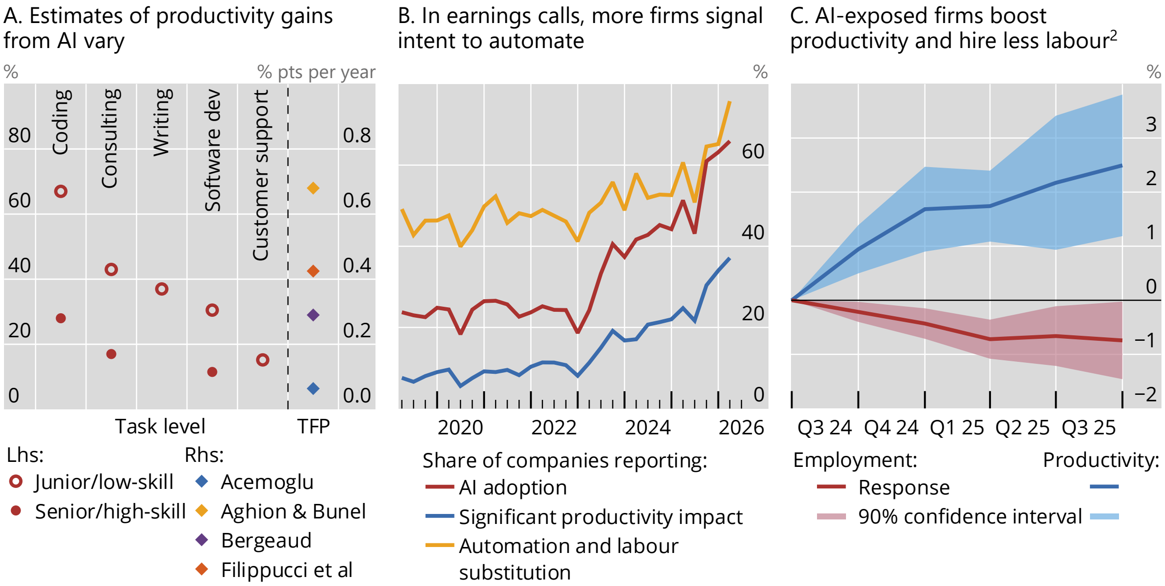

AI = artificial intelligence; software dev = software developers; TFP = total factor productivity.

[1] See additional notes to graphs for details. [2] Based on a cross-sectoral period-by-period regression where cumulative productivity (employment) growth is regressed on Q3 2024 log of productivity (employment) and Q3 2024 sectoral exposure to AI in the United States.

Sources: Acemoglu (2025); Aghion and Bunel (2024); Bergeaud (2024); Brynjolfsson et al (2025); Cui et al (2026); Dell'Acqua et al (2023); Filippucci et al (2024); Gambacorta et al (2024); Noy and Zhang (2023); S&P Global Market Intelligence; national data; BIS.

In the near term, the ongoing AI investment boom raises questions about the sustainability of the current economic expansion. The five largest hyperscalers are set to spend over a trillion US dollars on AI-related capital expenditure from 2025 through 2026. These commitments are outpacing earnings and the free cash flow of these firms, leading some to issue debt to raise additional financing (Graph 11.A). This investment race may be partly driven by the perception that only a small number of players with superior technology will ultimately dominate the market shares. The intense competition raises the risk of firms over-committing resources to investment projects with still uncertain returns, leaving all firms vulnerable to disappointments in AI payoffs. Model analysis based on such contest motives highlights the downside risk of current AI exuberance. As competitive pressure drives capex higher, the net economic surplus – the total payoff less investment costs – declines for the sector as a whole and could turn negative in adverse scenarios (Graph 11.B). Disappointment in returns could trigger a sudden pullback in financing and turn the capex boom into a protracted investment bust, with potential knock-on effects on financial conditions (see below).

Another risk is that the AI boom runs into a supply side roadblock. The AI build-out has recently been facing growing bottlenecks in electricity, advanced semiconductors and grid equipment. Fast-growing demand for computing power is already pressuring electricity prices and input costs, with potential spillovers to inflation. Looking ahead, these temporary shortages may also amplify over-investment, as firms attempt to lock in future capacity through long-dated contracts that further expose them to any disappointments in demand.

Historical episodes of investment booms offer instructive parallels (Graph 11.C). The canal mania of the 1830s, the British railway mania in the 1840s, the electrification

Rapid AI boom raises questions about its sustainability[1] Graph 11

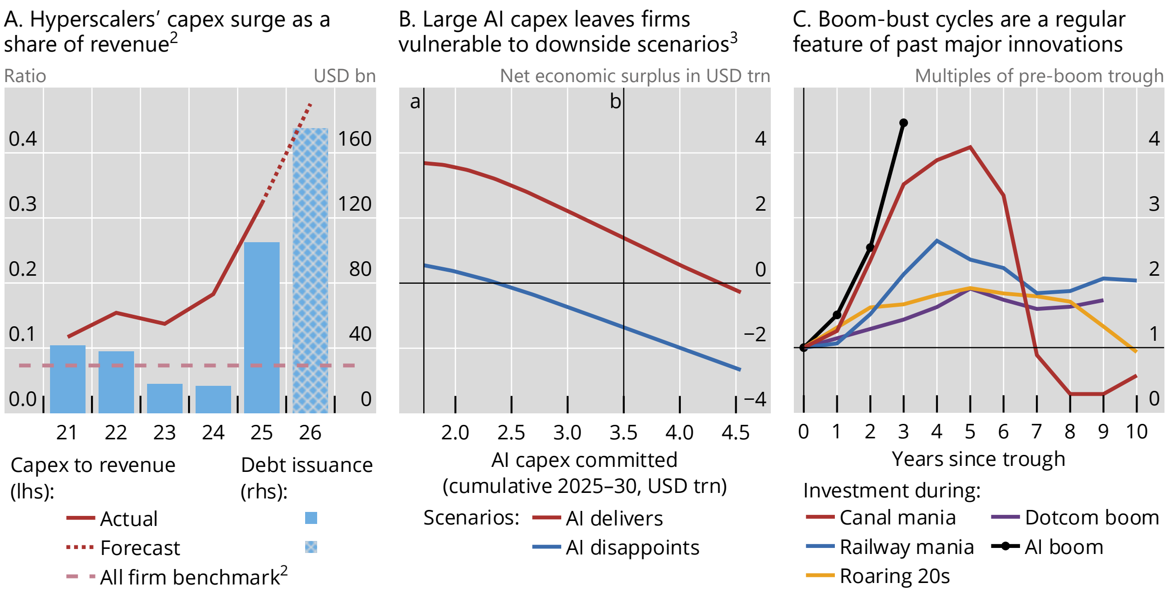

AI = artificial intelligence; capex = capital expenditure.

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

[1] See additional notes to graphs for details. [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 (2026b) for details.

Sources: Cranmer (1960); Rungcharoenkitkul (2026b); Bank of England; Federal Reserve Bank of St Louis; Bank of America; S&P Global Market Intelligence; companies' communications; national data; BIS.

exuberance of the late 1920s (roaring 20s) and the dotcom boom of the late 90s all shared one common trait: a genuine technological breakthrough that attracted capital in excess of what commercial returns could ultimately justify. These episodes ended with an eventual reversal in investment, inducing economy-wide recessions. The scale and pace of the current AI investment boom accompanied by expectations of large productivity payoffs bear resemblance to these precedents, highlighting potential downside risks in the near term.

Financial vulnerabilities as amplifiers

Should inflation rise significantly or AI-led investment turn to a bust, the macroeconomic consequences could be amplified by existing financial vulnerabilities. A tightening of policy rates needed to contain inflation could precipitate a sharp pullback in asset prices after a prolonged period of exuberant risk-taking, triggering disruptive macro-financial feedback loops. A reversal of AI optimism could likewise have major financial consequences, given AI firms' rising leverage and growing footprint in credit markets. Vulnerabilities extend to their supplier ecosystem, including engineering, procurement and construction (EPC) contractors whose balance sheets are comparatively weak, leaving them exposed to any capex pullback by hyperscalers.

Rich equity valuations are becoming a source of macroeconomic vulnerability — Graph 12

A. Stocks are pricing ambitious earnings growth expectations[1]

B. Risk compensation has fallen in US equity markets after the pandemic[1, 2]

C. US households increased exposure to stocks since dotcom bust

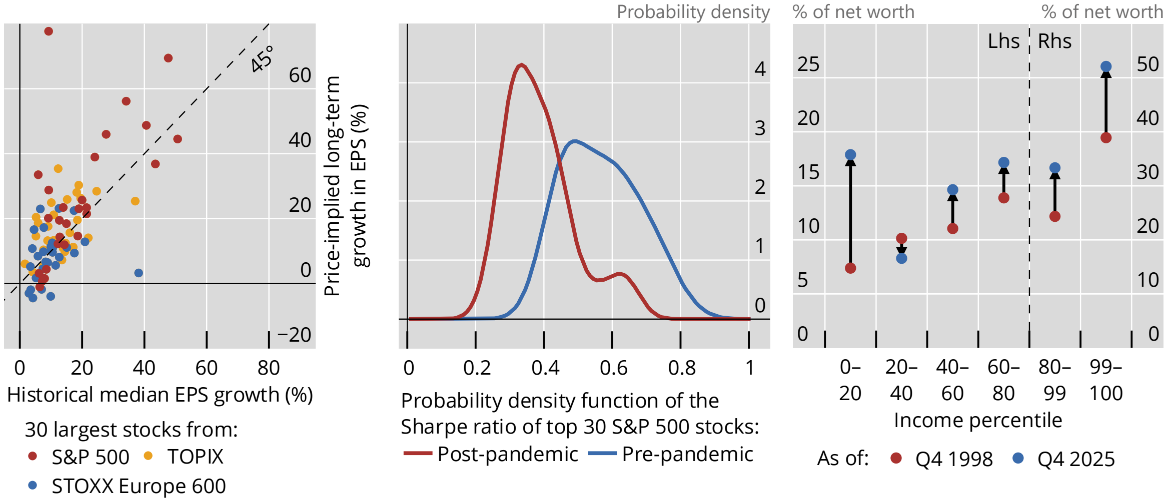

EPS = earnings per share; TOPIX = Tokyo Stock Price Index.

[1] See additional notes to graphs for details. [2] Sharpe ratio = excess return divided by volatility.

Sources: Board of Governors of the Federal Reserve System; Kenneth French Data Library; LSEG Datastream; LSEG Workspace; BIS.

Equity valuations are elevated, particularly for firms at the core of AI development. The implied long-term earnings growth for the largest corporations sits well above recent historical benchmarks (Graph 12.A), with US stocks often trading at large premia to peers in other major markets. These implied rates often exceed even the elevated growth that some of the technology firms have delivered in their relatively short lifetimes. As these firms mature and command a larger share of the market, sustaining such high growth could become increasingly challenging.

Sentiment has also been a major driver of current valuations. Risk premia on the largest US stocks have compressed markedly since the Covid-19 pandemic, with the distribution shifting clearly to the left (Graph 12.B). This points to growing investor complacency and reduced compensation for risk-bearing. Post-pandemic exuberance has been largely broad-based across sectors and countries, coinciding with the rapid rise of AI as an investment theme following the release of generative AI tools in late 2022.

A major equity market correction could have larger macroeconomic consequences today than in the past. Household equity exposures have grown over the past few decades, both relative to total wealth (Graph 12.C) and income (Graph 3.C). A large correction in valuations could have more pronounced wealth effects and sharper consumption pullback than in the past. And with US stocks accounting for an outsized share of global equity markets – about 64% of the MSCI Global index – the wealth impact from a US-led repricing could propagate globally.

Financial stability could also be at risk in the event of an AI bust. Fixed income markets are one obvious vulnerability, given the high volumes of debt issued by hyperscalers, AI labs and EPC firms (Graph 13.A). Should hyperscalers slow or halt the aggressive pace of capex deployment, many borrowers across the supply chain could struggle to replace lost revenue and service their debt. The credit spreads of some AI

Corporate credit is vulnerable to repricing on AI disappointments[1] Graph 13

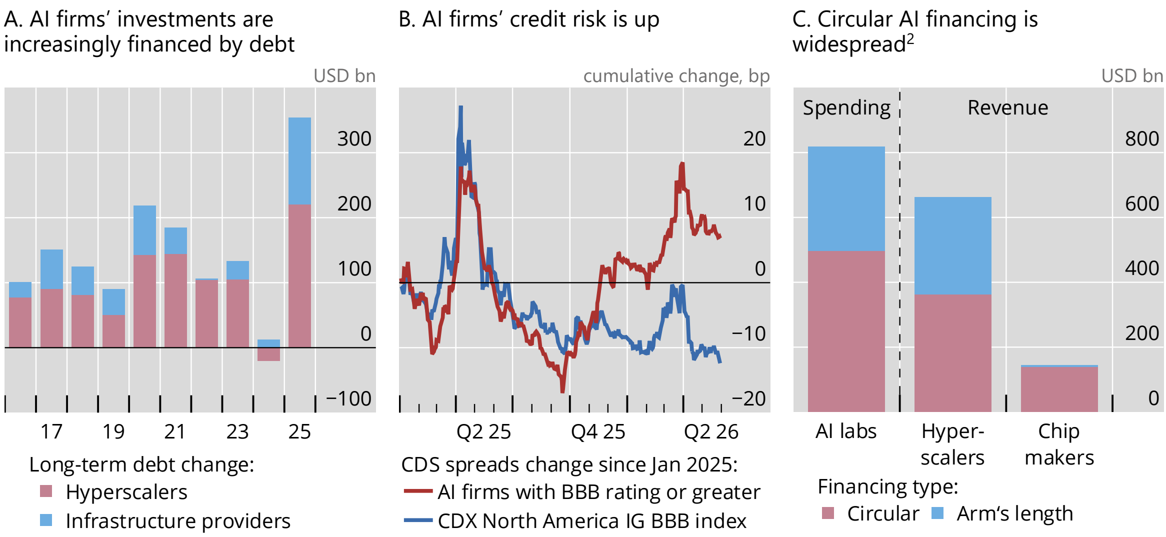

AI = artificial intelligence; CDS = credit default swap; IG = investment grade.

[1] See additional notes to graphs for details. [2] Announced multi-year commitments as of April 2026. Circular financing refers to a reciprocal investment structure in which hyperscalers take equity stakes in AI labs in exchange for the latter's purchase commitments, thus rechannelling capital back to investors as revenue; arm's length financing denotes flows that do not involve such circular financing.

Sources: Bloomberg; CNBC; LSEG Datastream; S&P Global Market Intelligence; The Wall Street Journal; BIS.

firms have already begun to widen somewhat to reflect this risk (Graph 13.B), even as equity markets continue to price in significant upside gains.

The opacity of AI-sector financing compounds these vulnerabilities. Hyperscalers, chip makers and AI labs are linked through a complex web of private arrangements. The most prominent is circular financing: chip makers and hyperscalers take equity stakes in AI labs or neocloud providers, who in turn commit to multi-year purchases of chips or computing power. Data centre construction is increasingly outsourced to third parties that lease facilities back to hyperscalers on long-dated contracts with embedded exit clauses. The terms of such deals are typically poorly disclosed, with risks of the same asset being pledged multiple times. Together, such arrangements account for a sizeable share of sector-wide financing and forward revenue (Graph 13.C).

A sharp repricing of equity risk could prompt a reassessment of corporate credit risk and lead to tighter credit conditions more broadly.[1] Indeed, broad indices of credit spreads tend to correlate negatively with stock market returns (Graph 14.A), more so for the high-yield than the investment grade segment. While large, synchronised corrections in both markets are rare, there are notable precedents such as the Great Financial Crisis and the March 2020 dash for cash episode. A repricing of risk this time, whether triggered by higher interest rates or an AI bust, has the potential to be similarly disruptive by triggering a corporate credit freeze with wider implications for aggregate investment.

Any tightening in credit conditions could expose existing vulnerabilities in the less transparent private credit space, whose reach has expanded among middle market and small firms (Graph 14.B). Signs of stress are already visible: direct lending funds catering to retail investors have faced mounting redemption requests, forcing some to liquidate assets and return capital despite having no contractual obligation to do so. A larger shock, whether from a renewed inflation surge or a sharp AI-led repricing, could trigger a more widespread credit crunch. The consequences could extend beyond the non-bank perimeter, given banks' growing and opaque exposure to private credit funds, compounded by overlapping ties through insurance companies' balance sheets. Given that the affected segment is smaller corporates that account for a large share of job creation, the real economy implications could be substantial.

The growing role of private credit also raises concentration risks. Direct lending funds, dominant players in the private credit ecosystem, have quadrupled their lending to the AI and information technology (IT) sectors in the past five years, to about 15% of their portfolios.[2] These loans tend to be larger than those in other sectors, while their terms such as tenor and pricing remain broadly similar, raising questions about lending standards and risk pricing. Investor enthusiasm has allowed more funds to participate, increasing concentration risks as software firms draw on multiple private credit lenders simultaneously (Graph 14.C).

Cyclical vulnerabilities are compounded by secular forces reshaping the financial system. NBFIs, beyond their roles in private credit activity, are shifting credit and leverage outside the banking sector and interacting with fiscal risks (Chapter II). Advances in AI and digitalisation are changing the contours of financial stability risk. Frontier AI models lower the cost and accelerate the pace of cyber attacks. Money-like digital assets such as stablecoins (Chapter III) introduce run-prone instruments outside the bank perimeter. These challenges are testing the adequacy of the current regulatory framework.

Sudden tightening of financial conditions may overly curtail credit supply — Graph 14

A. Corporate equity and credit risk are linked

B. Private credit has become a growing financing source to SMEs[1]

C. Exposure to software borrowers is concentrated[2]

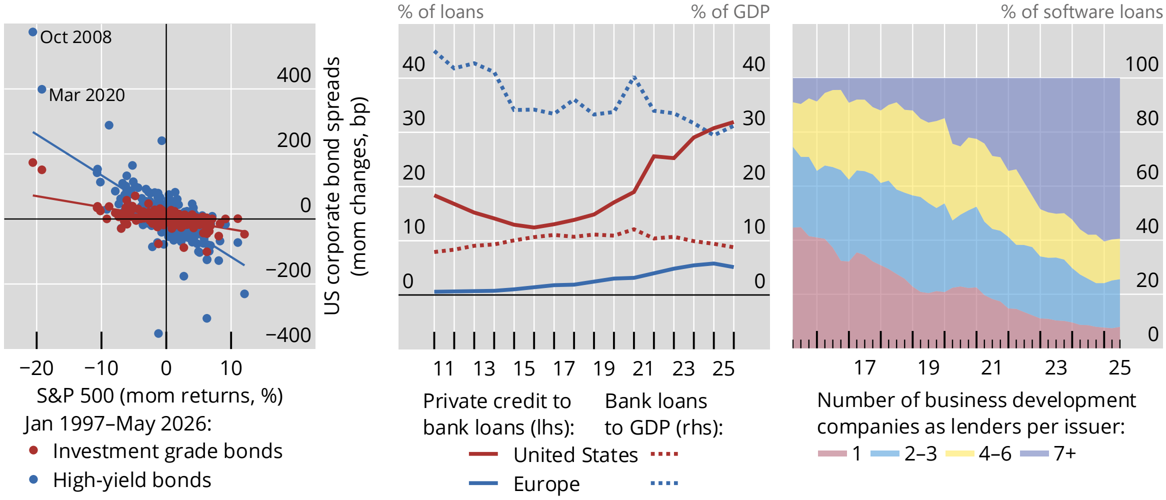

SMEs = small and medium-sized enterprises.

[1] See additional notes to graphs for details. [2] Number of distinct business development companies (BDCs) lending to the same borrower, expressed as a share of the total BDC loan volume to the software sector.

Sources: Avalos et al (forthcoming); ICE Data Indices; LSEG Datastream; Macrobond; PitchBook Data Inc; national data; BIS.