BIS

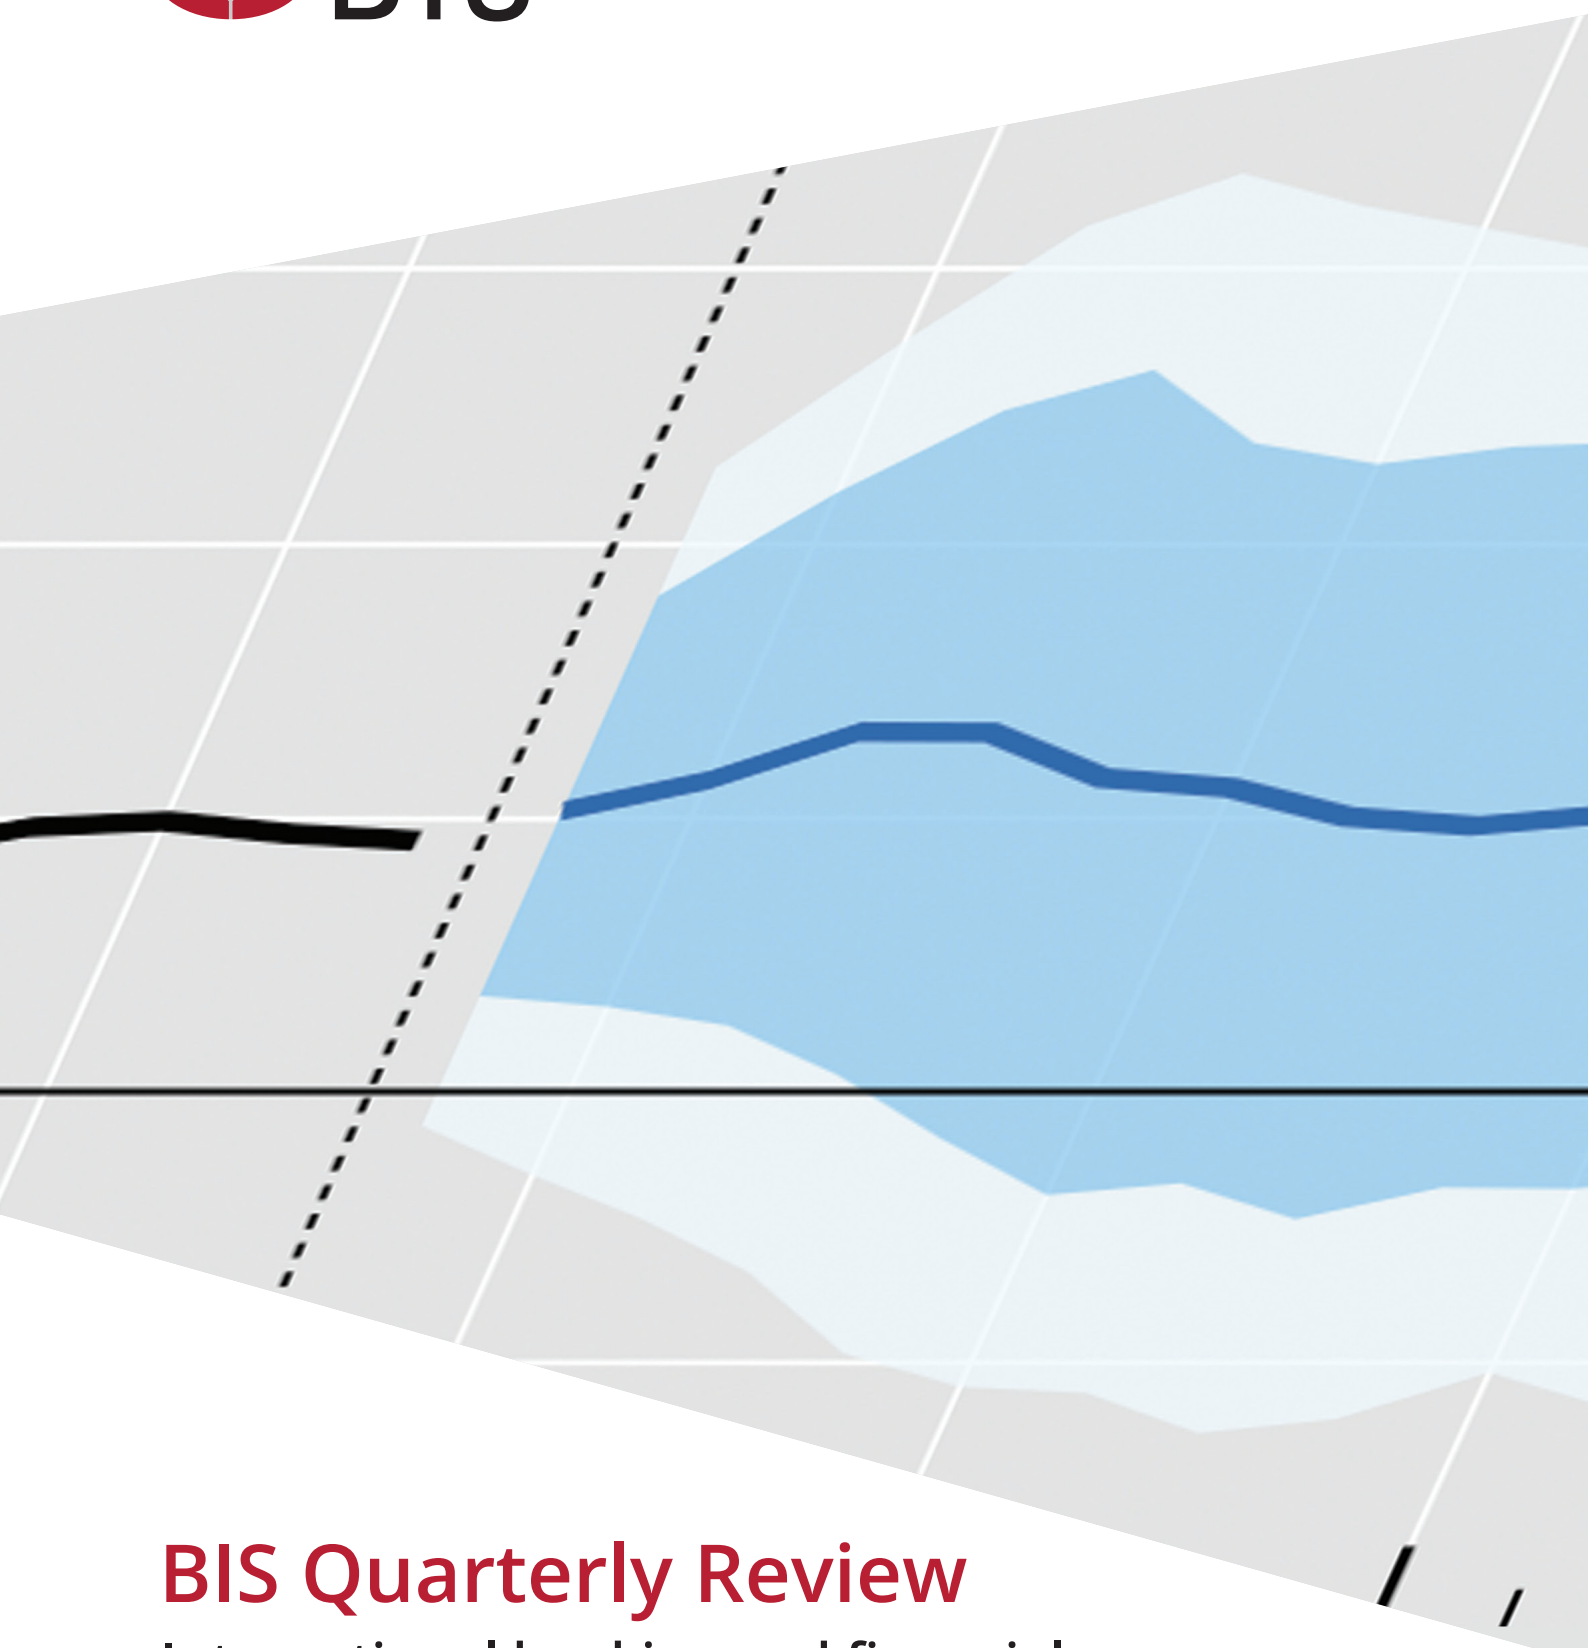

BIS Quarterly Review

International banking and financial market developments

March 2026

BIS Quarterly Review
Monetary and Economic Department

General queries concerning this commentary should be addressed to Andreas Schrimpf (tel +41 61 280 8942, e-mail: andreas.schrimpf@bis.org), queries concerning specific parts to the authors, whose details appear at the head of each section.

This publication is available on the BIS website (https://www.bis.org/publ/qtrpdf/r_qt2603.htm).

ISBN 978-92-9259-934-8 (online)
ISSN 1683-0121 (online)

BIS Quarterly Review

March 2026

International banking and financial market developments

Articles

Notations used in this Review

billion	thousand million
e	estimated
lhs, rhs	left-hand scale, right-hand scale
$	US dollar unless specified otherwise
…	not available
.	not applicable
–	nil or negligible

Differences in totals are due to rounding.

The term "country" as used in this publication also covers territorial entities that are not states as understood by international law and practice but for which data are separately and independently maintained.

Markets recalibrate amid shifting currents[1]

During the review period, financial markets had to adjust to shifting currents. Even though markets initially appeared calm on the surface, there were significant shifts under the surface as investors rotated away from previously high-performing assets. Volatility began to creep up, exacerbated by the conflict in the Middle East in early March.

Global equity markets saw regional and sectoral shifts in late 2025, with investors moving away from US large cap and growth stocks. European and Japanese equities posted gains, while emerging market economy (EME) equities saw an even greater boost. Despite solid earnings, concerns over artificial intelligence (AI) spending and disruption weighed on richly valued technology companies. Value and small cap stocks outperformed. While index volatility rose, it was eclipsed by significantly greater individual stock volatility.

Foreign exchange (FX) market shifts signalled growing investor unease in an increasingly fragile risk landscape. Through most of the review period, the US dollar depreciated as is typical of risk-on environments, but at the same time the Swiss franc appreciated notably, a hallmark of global risk-off episodes. The escalation of geopolitical tensions in the Middle East reversed the depreciation of the dollar, consistent with both its positive correlation with oil prices in recent years and the intensification of the risk-off sentiment.

Precious metals and energy prices also experienced notable volatility. Gold and silver surged early in 2026, reflecting a unique mix of investors' quest for safe havens and speculative interest. However, the rally ended abruptly in late January, in part due to leveraged position unwinds. Rising geopolitical tensions in early March drove oil and natural gas prices significantly higher. Amid these shifts, precious metals came under renewed pressure, trading more like risk assets than safe havens.

Credit markets remained fairly stable but showed some differences across segments. Spreads for investment grade and high-yield bonds stayed near historical norms, supported by resilient investor appetite for yield, despite impacts from escalating tensions. While issuance in bond markets broadly held up, activity in leveraged loans and private dealmaking declined from earlier peaks. Concerns over AI-driven disruption refocused attention on private credit portfolios, especially those with significant software exposures, as investor redemptions intensified.

Sovereign bond markets diverged across advanced economies (AEs). Long-term yields rose sharply in Japan, while those in the United States and the euro area moved sideways. After initial narrowing, euro area sovereign spreads widened again after the conflict broke out. Inflation expectations edged up, leading investors to revise expectations of policy rates upwards and push back the expected timing of US rate cuts. Bond yields moved up and yield curves mildly flattened as a result, reversing some of the previous steepening.

EME assets initially rose, driven by a weaker dollar, carry trades and portfolio inflows, but momentum waned with intensifying geopolitical frictions. Latin American and Europe, the Middle East and Africa (EMEA) currencies outperformed, while Asian currencies saw modest gains. Portfolio inflows fuelled equity rallies, especially in

[1] The review period covers 29 November 2025 to 5 March 2026.

Key takeaways

- *Risky asset returns diverged amid rotation shifts across regions and sectors, while elevated uncertainty and spikes in energy prices on the back of rising geopolitical tensions triggered bouts of volatility.*
- *Sovereign yields moved unevenly across AEs, and investors adjusted monetary policy expectations in response to higher expected inflation amid heightened geopolitical strains.*
- *Asset markets in EMEs initially rallied, supported by a softer US dollar and renewed investor appetite for carry trades, but momentum waned with escalating geopolitical tensions.*

commodity-exporting EMEs, and debt markets reflected favourable financing conditions. However, this positive trend was disrupted in early March due to heightened geopolitical risks.

Cross-currents: selective risk-taking collides with geopolitical tensions

Despite geopolitical flare-ups, investors' risk appetite broadly remained resilient through most of the review period, but risk-taking was selective. Solid corporate fundamentals and macroeconomic data releases contributed to this resilience. Investors' tolerance for risk was put to the test as a resurgence of tariff-related uncertainty, sharp moves in precious metal prices and rising geopolitical tensions triggered bouts of volatility. Concerns about high valuations called into question the sustainability of the momentum in technology stocks and led to underperformance as investors rotated to other sectors. Yet risky assets held their ground for the most part, with equity markets posting modest gains and credit spreads remaining compressed.

The equity market resilience that characterised the second half of 2025 was tested on several occasions during the review period, but broadly carried over into 2026. Equity indices rose, supported by solid risk appetite and fundamentals (Graph 1.A). The rise in US equity markets was relatively subdued, while European and Japanese equities posted stronger gains, reflecting improved sentiment towards their respective economies' outlooks and diversification flows away from US stocks. Rising tensions in the Middle East reversed some of the gains in early March, in particular for European and Asian equities, given their greater exposure to energy-related supply disruptions.

Large US technology stocks faced volatility despite strong earnings, as concerns over elevated valuations and future capital spending emerged. Guidance on higher capital spending raised fears of potential earnings disappointments, particularly for AI "hyperscalers" (see Box A). While corrections lowered valuations, price/earnings (P/E) ratios for the broader market remained above historical norms and approached dotcom bubble levels, though big tech firms' ratios were comparatively lower (Graph 1.B). As the Magnificent 7 (M7) stocks had driven their share of the S&P 500 index to nearly 35%, declines in these stocks began to drag on market indices.

Signs of regional rotation emerged as investors diversified away from concentrated US large cap exposure, with concerns over the scale of AI-related capital expenditure denting sentiment. Equity fund flows appeared to shift from the United States towards EMEs and other AEs, particularly those in Europe (Graph 1.C).

US equities lagged those of other AEs and EMEs as valuations stayed elevated[1] Graph 1

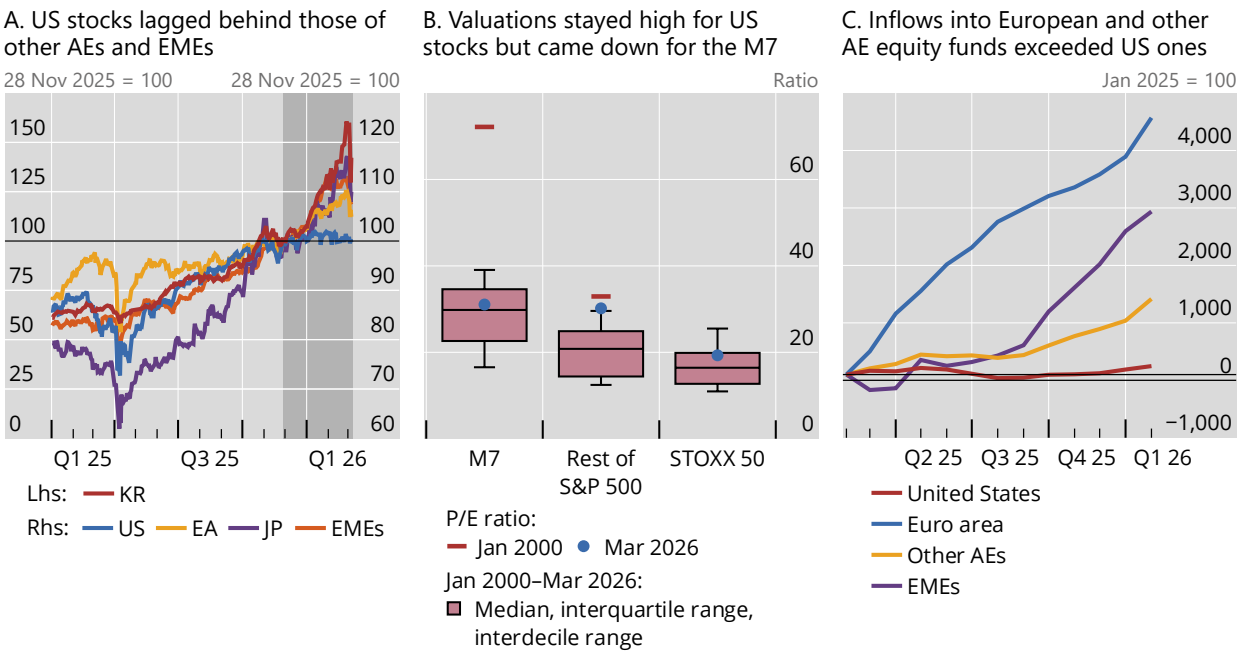

M7 = Magnificent 7; P/E = price/earnings.

The shaded area indicates 29 November 2025–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: IMF; Bloomberg; EPFR; LSEG Datastream; LSEG Workspace; BIS.

Box A

Financing the AI infrastructure boom: on- and off-balance sheet borrowing

*Egemen Eren, Ingomar Krohn, Karamfil Todorov*①

Investment in artificial intelligence (AI) infrastructure, particularly data centres, has risen rapidly and now accounts for a substantial share of investment in advanced economies. Against this backdrop, big US tech firms – "hyperscalers"② – have accelerated their capital expenditures (capex) (Graph A1.A). But, as these expenditures have grown significantly beyond their usual investments, a rising share of spending is funded through borrowing.③

Corporate bond markets have been hyperscalers' primary source of financing. Their gross issuance increased markedly, topping $100 billion in 2025 (Graph A1.B). Most issuance was long-term, with maturities over five years, locking in funding for multi-year build-outs. However, credit default swap (CDS) spreads rose (Graph A1.C), especially for hyperscalers with lower credit ratings, reflecting both the volume of supply and uncertainties around the projects' payoffs.

Alongside traditional bonds, hyperscalers have turned to off-balance sheet arrangements to finance infrastructure expansions, often in partnership with private credit firms. A common structure involves a dedicated vehicle – often a joint venture or special purpose entity – that acquires or develops data centre assets. The vehicle is capitalised with equity from a consortium of sponsors and raises debt through private placements. The hyperscaler typically holds a minority stake, commits to long-term operating leases or capacity offtake agreements (long-term commitments to purchase or reserve power and compute capacity), and may provide various guarantees. Economically, this substitutes upfront capex with multi-year operating expenses while keeping most of the associated debt off the hyperscaler's balance sheet. The debt is serviced by lease cash flows

and is held by private credit funds and other institutional investors, sometimes with investment grade features supported by asset backing and contractual guarantees from hyperscalers with investment grade credit ratings.

AI hyperscalers increasingly finance capital expenditures with debt[1]

Graph A1

A. Hyperscalers almost doubled capex over the past two years…

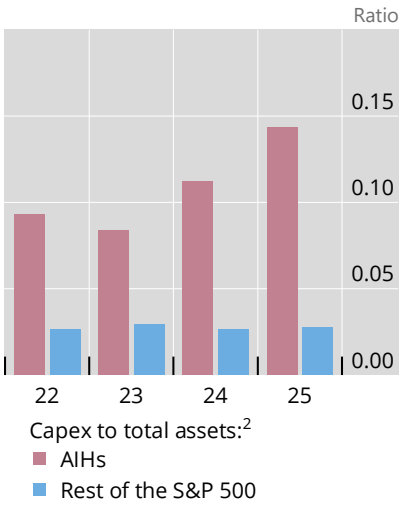

B. …and financed them by issuing large amounts of debt with long maturities[3]

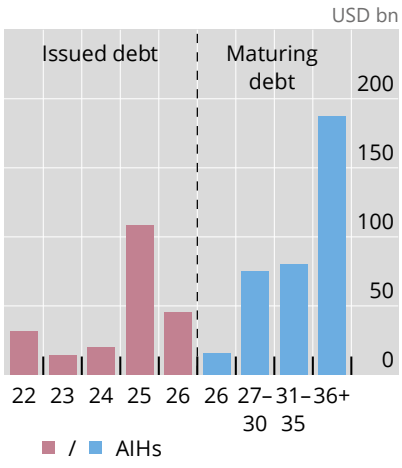

C. CDS spreads of AIHs rose compared with IG benchmarks

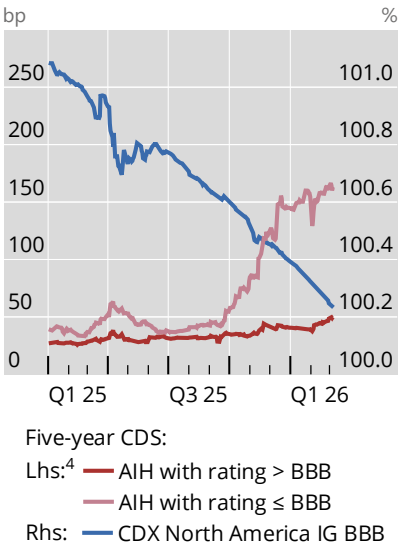

Capex = capital expenditure; CDS = credit default swap; IG = investment grade.

[1] AI hyperscalers (AIHs) = Alphabet, Amazon, Meta, Microsoft and Oracle. [2] Market capitalisation-weighted average. The composition of the S&P 500 is fixed on 5 March 2026. [3] Data as of 5 March 2026. [4] Simple average.

Sources: LSEG Datastream; LSEG Workspace; S&P Global Market Intelligence; authors' calculations.

These arrangements amount to "shadow borrowing": obligations that are economically akin to debt but largely reside outside corporate balance sheets. By channelling sizeable private credit into AI infrastructure, these structures strengthen links between hyperscalers and non-bank investors such as private credit vehicles and insurers. Banks support the vehicles with funding lines, potentially creating new shock transmission channels – eg via refinancing pressures at the vehicle level, procyclical shifts in private credit appetite or the activation of guarantees.

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks. ② These include firms such as Amazon, Alphabet, Microsoft, Meta and Oracle. ③ See also I Aldasoro, S Doerr and D Rees, "Financing the AI boom: from cash flows to debt", *BIS Bulletin*, no 120, 2026.

Supported by the dollar's depreciation, EMEs rallied for most of the review period, outperforming US equities (Graph 1.A, orange and blue lines). Asian AEs like Japan and Korea also delivered strong gains.

Within the United States, investors shifted from the M7 and high-momentum tech stocks to value and cyclical sectors, with small caps outperforming while software stocks faced pressure (Graph 2.A). Banks, energy, industrials, consumer staples and materials led gains (Graph 2.B), with banks benefiting from rising profitability in a resilient economy, and other sectors reflecting a broader rotation from growth to value stocks (Graph 2.C). Globally, software equities struggled due to concerns over AI-driven revenue disruptions, which also affected private credit firms with exposure to the sector (see Box B).

Overall risk appetite was sustained amid sectoral rotation[1]

Graph 2

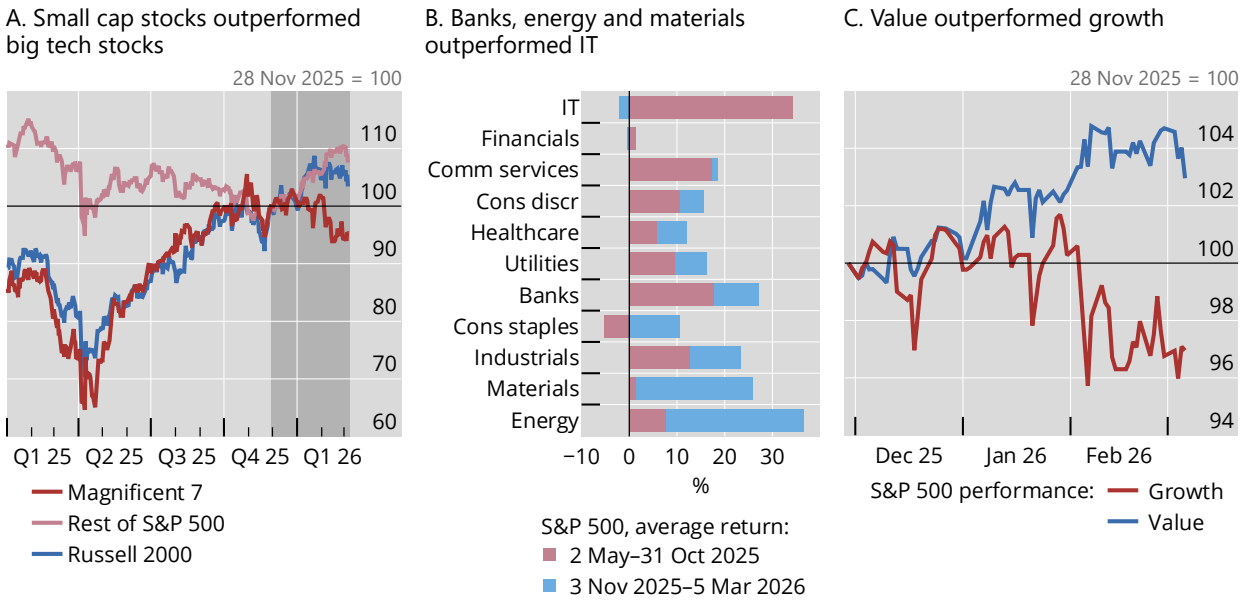

Comm = communication; cons = consumer; discr = discretionary; IT= information technology.

The shaded area indicates 29 November 2025–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: Bloomberg; LSEG Datastream; LSEG Workspace; BIS.

Index volatility crept up, reflecting sustained investor demand for downside protection, but it masked the even sharper rise in single-stock volatility under the hood. As volatility trended higher, geopolitical flare-ups brought about spikes to the VIX, but these were relatively short-lived. The dispersion index, which measures the gap between average single-stock and index volatility, climbed (Graph 3.A, red line). This signalled higher idiosyncratic risk relative to the uptick in index-level risk.

Currency movements diverged from their usual trends. A depreciating US dollar through most of the period was consistent with sustained risk appetite amid the rotation shifts. At the same time, appreciation pressures on the Swiss franc, which typically coincide with global risk-off episodes, underscored that investors balanced risk exposure with a demand for safety (Graph 3.B). The dollar depreciation was reversed after the outbreak of the Middle East conflict, reflecting the intensification of the risk-off sentiment after the conflict began and the dollar's positive correlation with oil prices in recent years.[2]

The conflict in the Middle East triggered turbulent moves in oil and gas prices towards the end of the review period. Prices of oil and natural gas (especially in Europe) jumped (Graph 3.C, blue and purple lines) and oil volatility spiked (Graph 3.A, blue line), reflecting the significant uncertainty in energy markets. Oil and gas futures curves shifted upwards, with the largest increases concentrated in shorter-term contracts.

[2] On the changing correlations of oil prices and the US dollar, see B Hofmann, D Igan and D Rees, "The changing nexus between commodity prices and the dollar: causes and implications", *BIS Bulletin*, no 74, 2023.

Private credit's software lending meets AI disruption

*Sebastian Doerr, Egemen Eren, Ingomar Krohn, Karamfil Todorov*①

Lending by private credit funds to software-as-a-service (SaaS) firms has grown rapidly and makes up a substantial share of funds' loan portfolios. Outstanding loans to SaaS firms increased from almost $8 billion in 2015 to over $500 billion, or 19% of total direct loans, by end-2025 (Graph B1.A).② By now, a third of private credit funds have extended loans to the SaaS sector, on top of their rising exposure to big US tech firms and other artificial intelligence (AI) companies (see Box A). In this box, we zoom in on business development companies (BDCs) to get a consistent picture of private credit lending to SaaS firms. BDCs are publicly traded and have quarterly disclosures. As such, they provide a window into the opaque private credit space. Moreover, BDCs account for one fifth of all direct loans in the United States and extended over 15% of their loans to SaaS firms in 2025.

Concerns that AI may disrupt traditional SaaS business models have led to notable price adjustments in the software sector. Software companies' stocks collapsed by almost 30% between October 2025 and February 2026 (Graph B1.B, red line), while BDCs' stock prices fell by about 10% on average (blue line). Meanwhile discounts to net asset value, which is largely determined by the book value of illiquid private loans, deepened (yellow line), potentially signalling worries about underlying valuations.

Worries about software stocks spilled over to private credit

Graph B1

A. Direct lending to SaaS firms is growing fast[1]

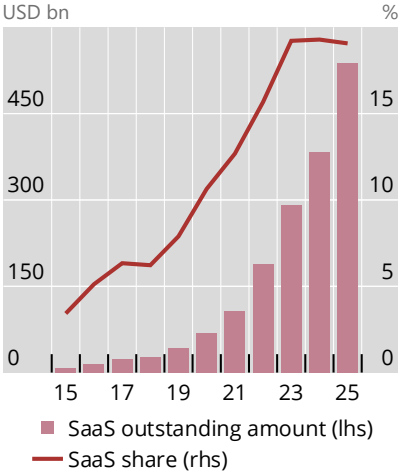

B. Stocks of software firms and BDCs fell in tandem

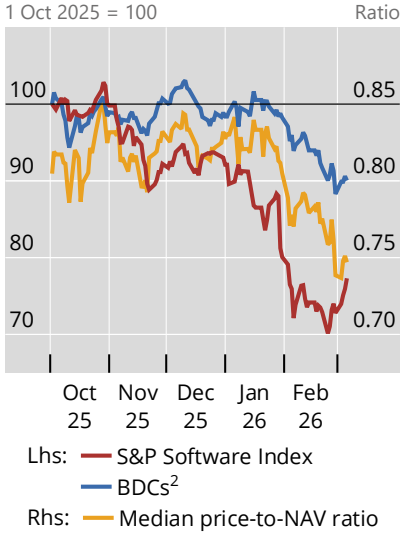

C. BDCs with higher SaaS exposure saw steeper declines[3]

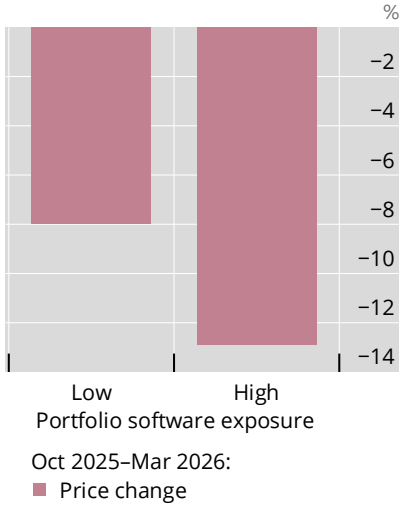

BDCs = business development companies; NAV = net asset value; SaaS = software as a service.

[1] Total outstanding direct loans to firms operating in the Pitchbook-defined vertical "SaaS". [2] Simple average of stock prices of 43 BDCs. [3] The sample comprises 37 BDCs. Price changes are calculated as the simple average of stock price changes within each subgroup between 1 October 2025 and 5 March 2026. The sample is split into groups with above- and below-median software exposure as of Q3 2025.

Sources: US Securities and Exchange Commission; Bloomberg; LSEG Datastream; PitchBook Data Inc; authors' calculations.

BDCs with higher SaaS exposure have underperformed their peers (Graph B1.C). Since October 2025, BDCs with high exposure to software firms have performed around 5 percentage points worse than those with low exposure. These developments highlight investor concerns that further advances in AI tools may disrupt the SaaS sector amidst redemption pressures from private credit's push towards retail investors.③

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks. ② Globally, private credit stands at over $2 trillion; see F Avalos, S Doerr and G Pinter, "The global drivers of private credit", *BIS Quarterly Review*, March 2025. ③ See I Aldasoro, S Doerr and K Todorov, "Retail investors in private credit", *BIS Bulletin*, no 106, 2025.

Precious metals posted gains over the period but were subject to large price swings. Gold and silver extended their 2025 momentum, reflecting a unique mix of safe haven demand and amplification by trend-following investors and leverage (see Box C).[3] Gold reached fresh highs in January 2026, and silver benefited additionally from industrial demand and supply shortages (Graph 3.C). The price of silver, which had doubled in 2025, increased over 50% in January 2026 alone. However, a sharp drawdown occurred in late January as gold and silver sold off rapidly amid the resolution of uncertainty surrounding the future chair of the Federal Reserve. Silver recorded its largest one-day loss since the 1980s, falling by nearly 30%. Precious metal prices extended their decline despite the escalation of the Middle East conflict, a situation that would typically boost demand for safe haven assets.

Crypto saw large sell-offs as part of the general rotation away from growth assets, but also an upswing amid the escalation of geopolitical tensions. Following the drop in technology stocks, bitcoin slumped by about 50% from its 2025 highs and touched 2024 levels (Graph 3.C). These moves were probably exacerbated by liquidations of leveraged long crypto positions.[4] Bitcoin's behaviour amid the escalating tensions early March was puzzling, but overall it ended the period above pre-conflict levels, pausing its prior decline.

Volatility increased, especially in gas and oil, while safe haven demands grew[1]

Graph 3

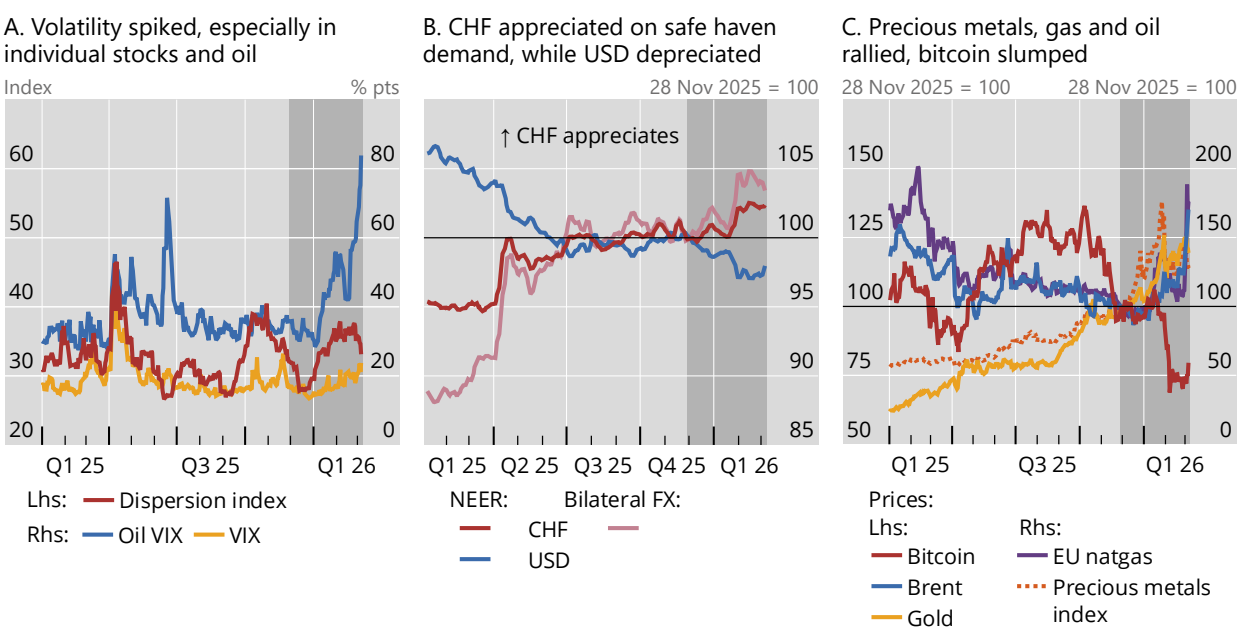

NEER = nominal effective exchange rate.

The shaded area indicates 29 November 2025–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: Bloomberg; CoinDesk Data; national data; BIS.

[3] See also G Cornelli, M Lombardi and A Schrimpf, "Bubble conditions in US equities and gold?", *BIS Quarterly Review*, December 2025.

[4] On the role of risk amplification due to leverage in crypto markets, see also M Schmeling, A Schrimpf and K Todorov, "Crypto carry", *Management Science*, forthcoming.

Box C

Boom and bust of the recent silver and gold rush: the role of leveraged retail investors

*Egemen Eren, Ingomar Krohn, Karamfil Todorov*①

After a prolonged rally through 2025 and into early 2026, prices of precious metals such as gold and silver reversed abruptly in late January and February 2026. Retail-driven exuberance, increasingly channelled through exchange-traded funds (ETFs), set the stage for outsize moves, continuing the trend from 2025.② The daily rebalancing of leveraged ETFs and margin-triggered liquidations amplified the swings, particularly in silver.

Following substantial gains in 2025 and a further surge in early January 2026, gold and particularly silver prices plunged in late January (Graph C1.A). After doubling over 2025 and rising by over 50% in January 2026, the silver price fell by about 30% in a single day in late January (blue line). Gold broadly followed a similar but less extreme pattern (red line). The precious metals crash seemingly coincided with shifts in expectations about the US dollar and the path of monetary policy, but was hard to square with broader changes in fundamentals. The abrupt price drop and the spike in precious metals' volatility point to the role of retail flows, and amplification of price moves due to forced sales by leveraged ETFs, trend-following investors such as commodity trading advisers (CTAs) and margin dynamics.

Precious metals rally was fuelled by speculative retail flows

Graph C1

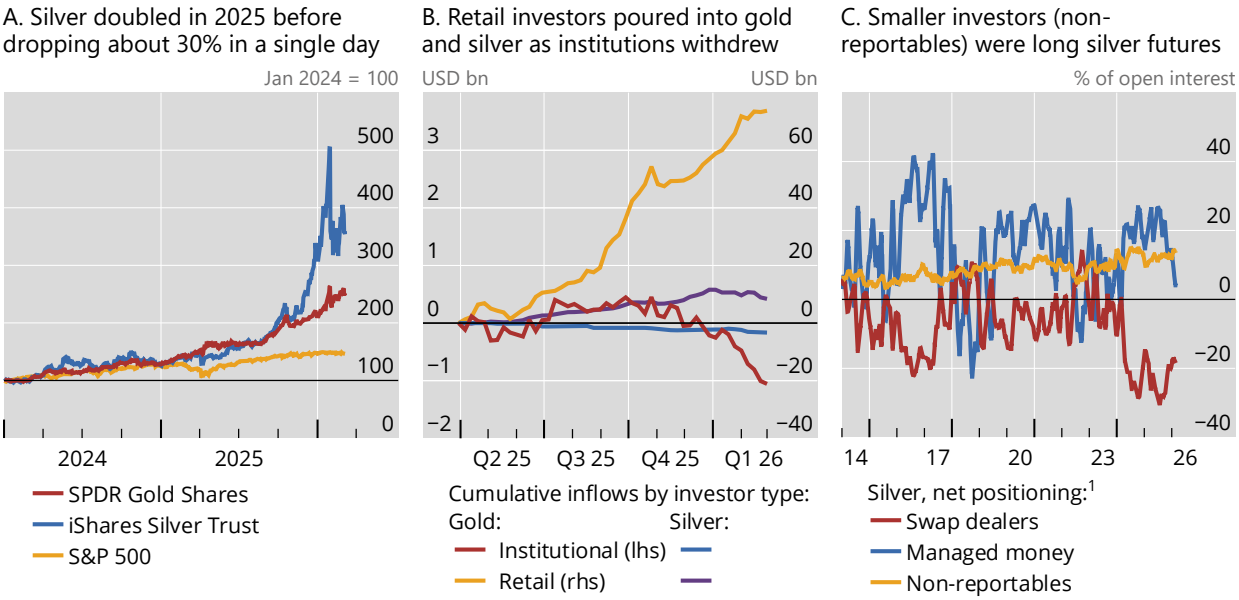

[1] Twenty-one-day moving average.

Sources: Commodity Futures Trading Commission; Bloomberg; LSEG Lipper; authors' calculations.

Fund flow data indicate that retail investors were the main source of inflows into silver and gold funds in the run-up to the episode. In contrast, institutional investors maintained stable positions or even trimmed exposure (Graph C1.B). Moreover, futures positioning reveals long leveraged exposure among smaller speculative participants. "Non-reportables" – typically smaller investors – were long silver futures heading into the correction. As prices fell sharply and exchanges raised margin requirements, these investors probably had to reduce positions quickly. "Managed money" – including CTAs and institutional investors – also cut long positions, while dealers stepped in to provide liquidity by cutting short positions (Graph C1.C).

For some time, retail investors have found it attractive to use ETFs to obtain precious metal exposure. Sustained premia of gold and silver ETFs over their net asset value (NAV) signalled strong, one-sided buying pressure that outpaced primary market arbitrage (Graph C2.A). Such persistent premia arise when demand for ETF shares exceeds the capacity of authorised participants to create new shares and deliver physical metal to

bring market prices down to NAV. As prices reversed in late January, these premia compressed rapidly, and for silver turned into pronounced discounts, consistent with one-sided selling pressure and a sharp turn in flows.

Leveraged ETFs rebalancing and margin increases amplified silver volatility

Graph C2

A. Persistent ETF premiums signalled consistent buying pressure

B. Leveraged ETFs' trading amplified volatility…[2]

C. …as margins increased

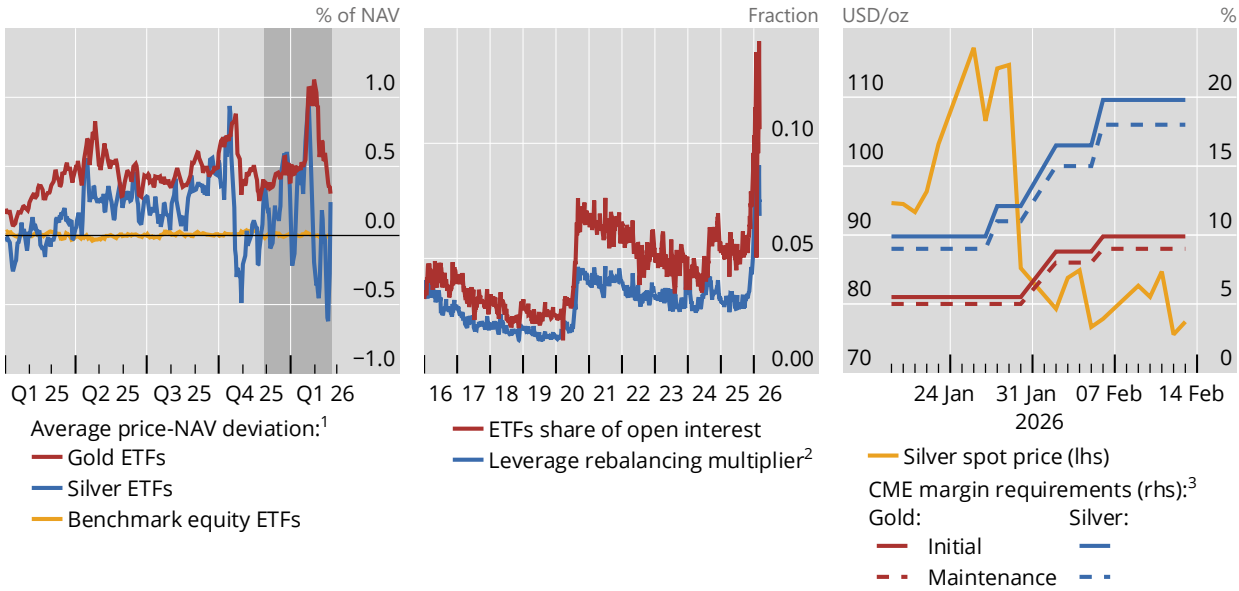

The shaded area indicates 29 November–5 March 2026 (period under review).

[1] Data winsorised at the 0.5th and 99.5th percentiles. Five-business day moving averages. For benchmark equity ETFs, simple average of iShares Core S&P 500 ETF, SPDR S&P 500 ETF Trust and Vanguard 500 Index Fund ETF. For Gold ETFs, simple average of 205 ETFs. For silver ETFs, simple average of 60 ETFs. [2] The leverage rebalancing multiplier shows what fraction of the futures market would need to be sold by leveraged ETFs when the market goes down. For example, a leverage rebalancing multiplier of 0.1 means that if silver goes down by 30% (as in late January), leveraged ETFs would need to sell 3% of the futures market. For more details, see K Todorov, "When passive funds affect prices: evidence from volatility and commodity ETFs", *Review of Finance*, vol 28, no 3, May 2024. [3] For COMEX 100 Gold Futures and COMEX 5000 Silver Futures, respectively.

Sources: Bloomberg; CME; LSEG Lipper; authors' calculations.

Leveraged silver ETFs contributed to the turmoil because of their amplification mechanics. To maintain fixed daily leverage, these funds rebalance each day. When prices rise, they buy the underlying asset (often via silver futures) to restore target leverage and when prices fall, they sell the underlying asset. This predictable, momentum-like trading creates feedback loops that reinforce prevailing trends and can distort prices.③

The footprint of leveraged ETFs' destabilising trading appears to have grown amid the retail-driven exuberance in precious metal markets. The leverage rebalancing multiplier, a summary measure of the market impact of leveraged ETFs' daily rebalancing flows, doubled over the course of 2025 (Graph C2.B, blue line). The share of ETFs in the market followed similar dynamics (red line). This indicates that leveraged ETF activity became a larger part of the market and intensified price trends.

Margin-triggered liquidations further amplified the sell-off. Rapid price declines increased variation margins on futures positions, and several exchanges tightened initial margin requirements during the episode (Graph C2.C). The resulting funding pressures forced deleveraging among participants most exposed to the downdraft, akin to past stress episodes.④ The liquidations of investors' positions, alongside systematic selling from leveraged ETF rebalancing into the decline, probably added to downward pressure, creating a self-reinforcing loop of lower prices and further margin calls.

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks. ② G Cornelli, M Lombardi, and A Schrimpf, "Bubble conditions in US equities and gold?", *BIS Quarterly Review*, December 2025. ③ K Todorov, "When passive funds affect prices: evidence from volatility and commodity ETFs", *Review of Finance*, vol 28, no 3, May 2024. ④ BIS, "Carry off, carry on", *BIS Quarterly Review*, September 2024.

Credit spreads remained compressed compared with their historical norms in the United States and Europe, especially in the high-yield segment (Graph 4.A). Resilient macroeconomic fundamentals and solid risk appetite helped to sustain valuations in these markets in spite of negative impacts from escalating tensions in the Middle East.

Despite compressed spreads and healthy investor demand, credit markets were marked by sluggish bond issuance and renewed focus on private credit portfolios. Investment-grade bond issuance remained relatively subdued despite usual seasonal upswings (Graph 4.B). Leveraged loan spreads broke their downward trend and picked up notably towards the end of the review period, signalling a negative shift in market sentiment (Graph 4.C). Leveraged loan issuance and private deal-making experienced significant declines from their previous peaks, highlighting a broader slowdown in riskier segments of the credit market (Graph 4.B, red and yellow lines). Strains emerged in private credit as investor redemptions intensified and some funds imposed redemption restrictions.

One notable development in credit markets was the increased borrowing of large technology firms. These firms, known as AI hyperscalers, provide the infrastructure, storage and specialised hardware necessary to develop and run large-scale AI models. To finance their massive infrastructure investments, these firms have increasingly raised funds through corporate debt and other complex funding structures. Market concerns about the sustainability of such borrowing were reflected in the widening of CDS spreads for some AI hyperscalers (see Box A).

Credit market issuance was lacklustre, and spreads diverged across segments

Graph 4

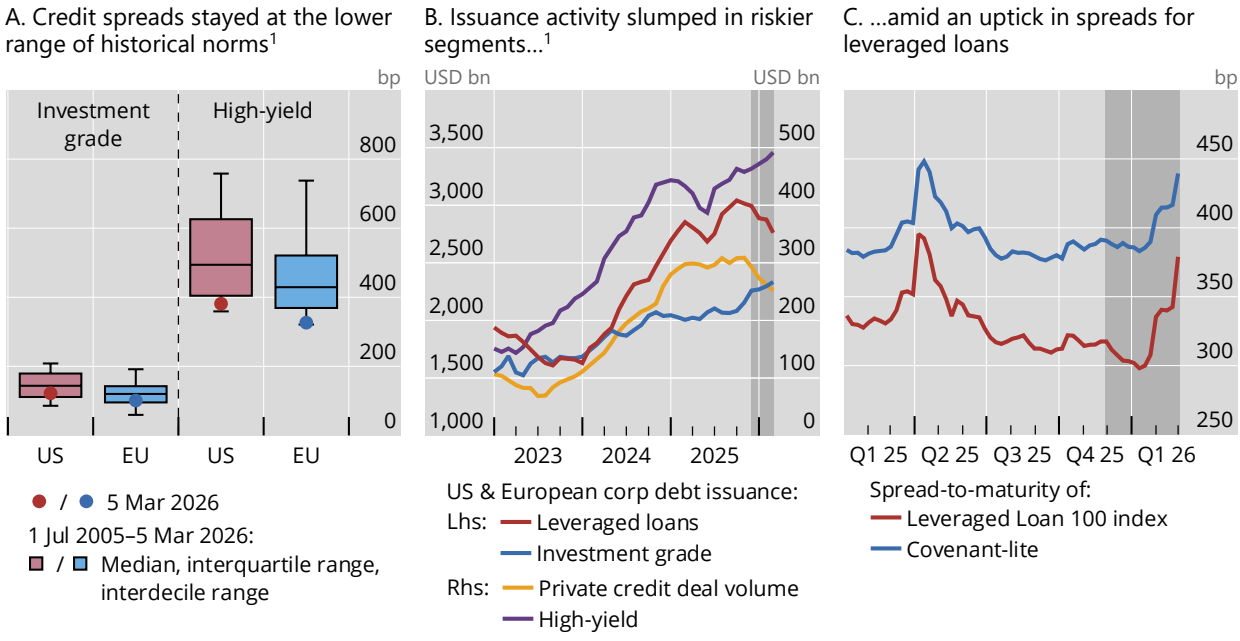

The shaded area indicates 29 November–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: Bloomberg; Dealogic; ICE Data Indices; PitchBook Data Inc; PitchBook Data Inc | LCD; BIS.

Sovereign bond markets diverged amid changing monetary and fiscal policy outlook

Sovereign bond markets evolved unevenly across AEs. Long-term yields on Japanese and Australian government bonds increased markedly (Graph 5.A, yellow line and purple line), while they moved sideways in the United States and Germany (Graph 5.A, red and blue lines). The volatility of US Treasuries edged up during the review period, in particular after the start of the conflict in the Middle East, even though it remained subdued by historical standards (Graph 5.B). In the euro area, sovereign spreads over German government bonds reversed their narrowing trend alongside rising geopolitical tensions in the Middle East as increasing oil and gas prices created uncertainties around the fiscal outlook (Graph 5.C).

The Middle East conflict led to a repricing of monetary policy expectations. In the United States, expectations of rate cuts shifted further out into the year, while in other AEs market participants expected more rate hikes during 2026. The re-evaluation of the path of interest rates reflected higher near-term inflation expectations, especially since the start of the conflict (Graph 6.B). Bond yields moved up as a result, further compounded by the anticipated fiscal implications of the conflict with an expected increase in debt supply that would eventually need to be absorbed by markets.

Sovereign bond market volatility edged up amid geopolitical developments — Graph 5

A. Long-term yields in Japan and Australia rose markedly

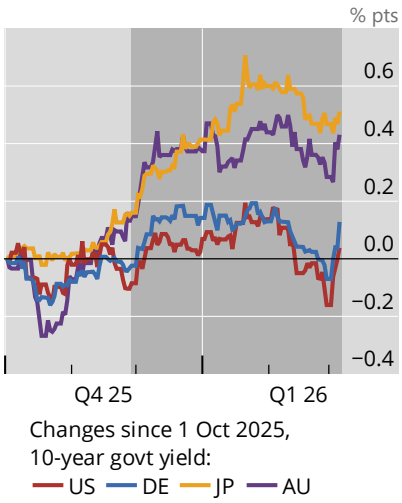

B. Treasury market volatility picked up in 2026

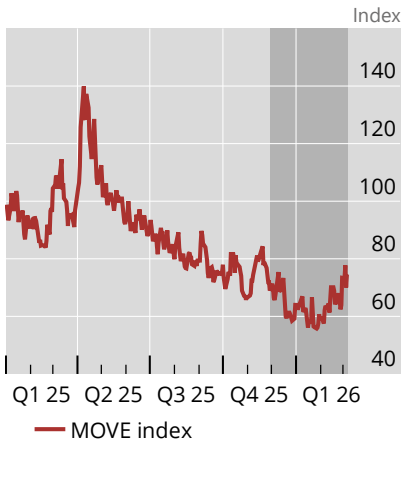

C. Yield differentials within the euro area widened following the start of the Middle East conflict

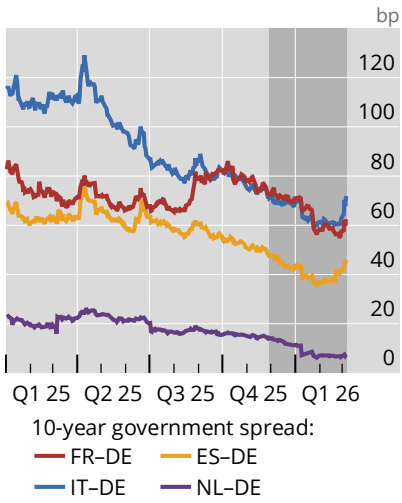

The shaded area indicates 29 November–5 March 2026 (period under review).

Sources: Bloomberg; BIS.

Policy rate and inflation expectations rose as geopolitical tensions escalated, while term spreads reversed[1]

Graph 6

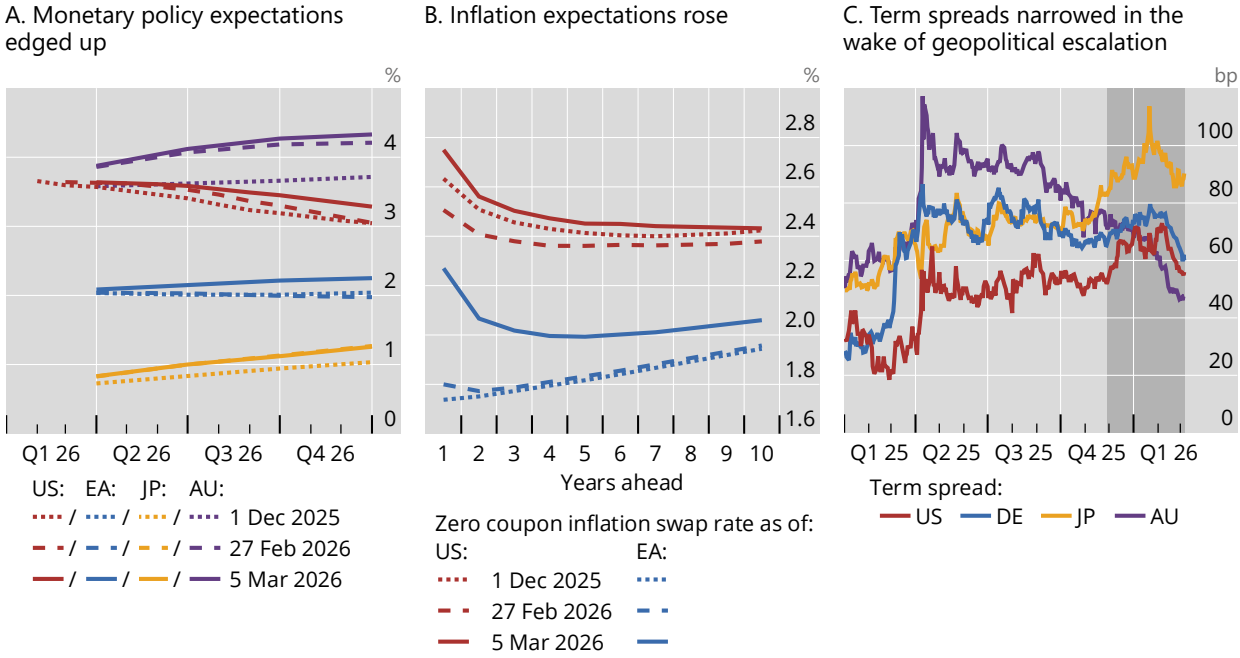

The shaded area indicates 29 November–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: Bloomberg; BIS.

Term spreads also moved differentially over time and across AEs. At the start of the review period, the overall tendency was for a steepening as long-term yields rose more than short-term ones, particularly in Japan. Expectations of large fiscal spending plans and lacklustre demand for long-term bonds by key players, such as life insurance companies, were key factors lifting Japanese long-term yields. Towards the end of the review period, the trajectory of term spreads reversed (Graph 6.C). This configuration primarily reflected a market pricing of higher near-term inflation expectations. In contrast to other AEs, yield curves flattened in Australia throughout the period, extending the earlier trend.

Global repositioning moves EME assets

EME assets initially benefited from the US dollar's weakness and portfolio rotation away from the United States, but the spike in uncertainty due to geopolitical strains later eroded some of the gains. Supportive market conditions towards the end of 2025 filtered through, with most EME currencies firming, bond spreads narrowing and equities advancing. Performance, however, varied across countries, reflecting differences in terms of trade, external vulnerabilities and domestic macroeconomic fundamentals.

EME currencies mostly appreciated against the US dollar throughout most of the review period (Graph 7.A), in line with the sustained initial risk appetite and rotation towards investments outside the United States. Currency gains were strongest among high-yielding, commodity-linked markets in Latin America and EMEA, while appreciation was more modest across Asia. In particular, the South African rand and Mexican peso posted notable increases, continuing the positive momentum built during 2025. These currencies were also supported by favourable interest rate differentials and broadly stable macroeconomic policy settings. With the escalation of the Middle East conflict, however, the resulting broad dollar appreciation halted the momentum across most EME currencies.

The conflict-related volatility in early March also affected carry trade profitability and disrupted the momentum dynamics which had lifted many EME currencies over previous months. While carry-to-risk ratios had picked up significantly in Mexican peso and Brazilian real vis-à-vis Japanese yen, they peaked in 2026, consistent with a rising risk of funding currency appreciation (Graph 7.B). In line with this, currency futures positioning indicated that, early on, leveraged funds increased their net long exposure to EME currencies – notably the Mexican peso – while at the same time keeping short positions in funding currencies such as the yen (Graph 7.C). These short positions were gradually pared back during the first few weeks of 2026 and were further scaled back when markets switched to risk-off in early March.

EME currencies' positive momentum was bolstered by carry trades, but came to a halt as uncertainty increased in early March

Graph 7

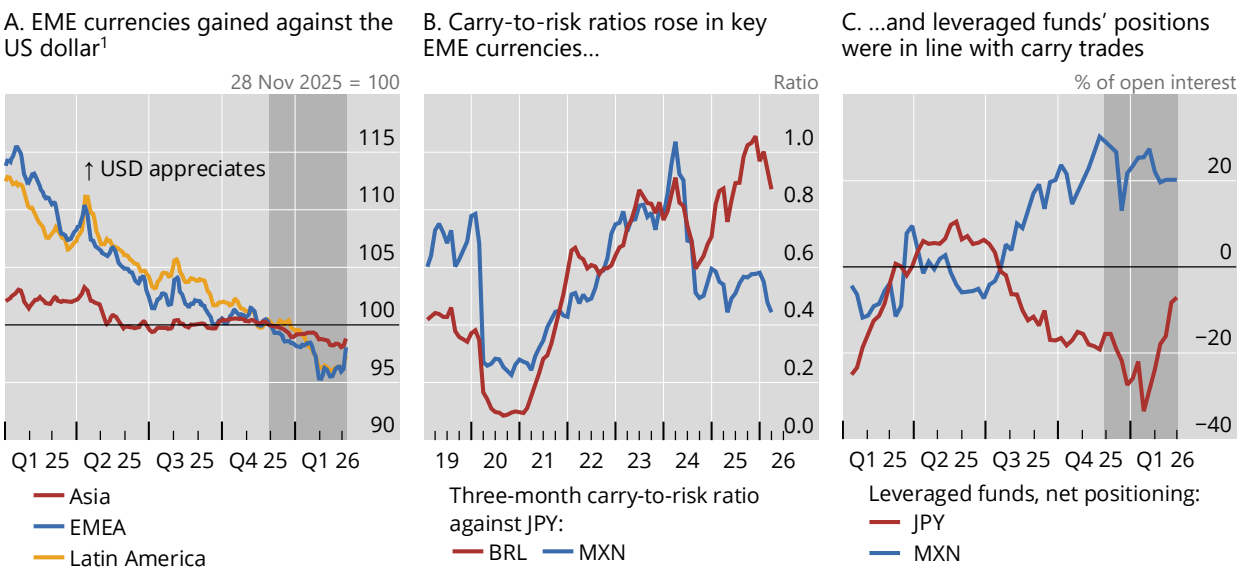

EMEA = Europe, the Middle East and Africa.

The shaded area indicates 29 November–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: Commodity Futures Trading Commission; Bloomberg; JPMorgan Chase; LSEG Datastream; national data; BIS.

Bond markets in EMEs broadly followed a similar pattern. Amid the rotation shifts away from US assets, investors extended their search for yield into EME debt. EME sovereign issuance gathered momentum as market conditions improved (Graph 8.A). The momentum faded at the start of 2026 as investment grade supply dwindled and high-yield volumes stayed flat. In parallel, spreads on US dollar-denominated EME sovereign bonds initially tightened but then reversed amid rising geopolitical tensions (black line).

Equity markets in EMEs mostly rallied early on but later reversed part of their gains. The initial momentum was widespread. Commodity-exporting economies in Latin America – particularly Peru, Brazil and Colombia – recorded gains, supported by rising commodity prices (Graph 8.B). Equity returns in EMEA were similarly strong, benefiting from diversified positioning across EMEs. In Asia, returns were lower but mostly positive, with some exceptions. Investor demand for EME assets was also evident in portfolio flows. Building on the pickup in inflows at the end of 2025, investors increased allocations to EMEs, particularly in EMEA and Latin America (Graph 8.C). Such momentum waned amid the conflict-related pressures, with stock markets of countries with large energy market dependencies falling more sharply than their peers.

Debt and equity markets in EMEs benefited from benign conditions, but momentum faded

Graph 8

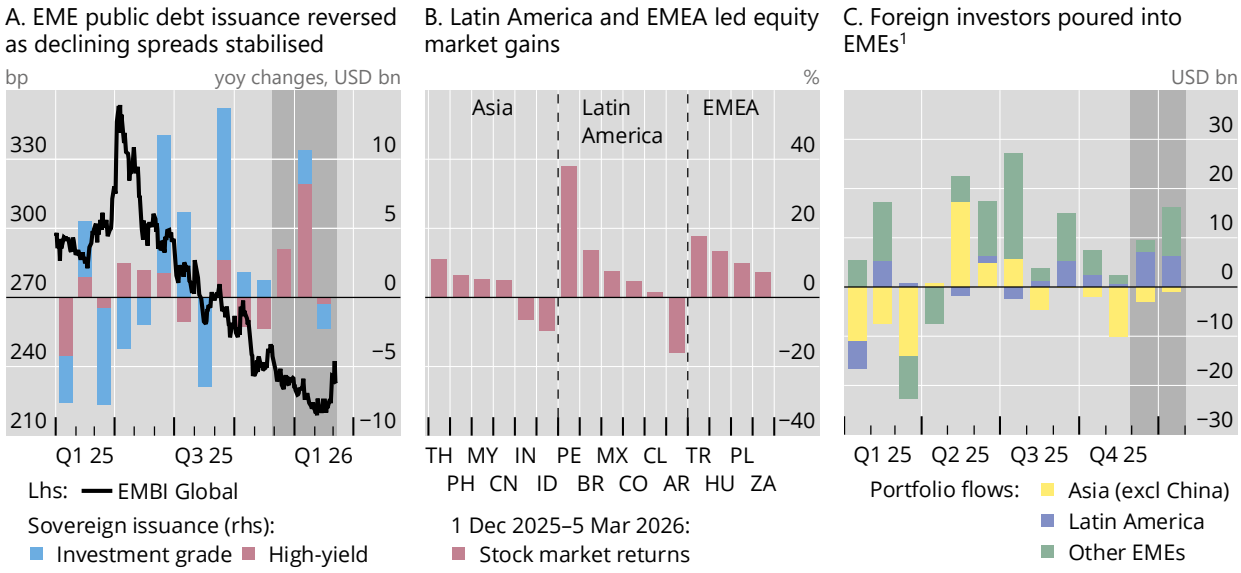

EMEA = Europe, the Middle East and Africa.

The shaded area indicates 29 November–5 March 2026 (period under review).

[1] See endnotes for details.

Sources: IIF; Bloomberg; Dealogic; LSEG Datastream; BIS.

Endnotes

Graph 1.A: EA = STOXX Europe 600; EMEs = simple average of stock market returns in AR, BR, CL, CN, CO, HU, ID, IN, MX, MY, PE, PH, PL, TH, TR and ZA; JP = Nikkei 225; KR = KOSPI; US = S&P 500.

Graph 1.B: Magnificent 7= Alphabet, Amazon, Apple, Meta, Microsoft, Nvidia and Tesla. The composition of the S&P 500 is fixed on 5 March 2026. Market capitalisation-weighted average. For STOXX 50, the series starts in June 2001. For March 2026, data as of 5 March 2026.

Graph 1.C: Cumulation since January 2025. EMEs = AR, BR, CL, CN, CO, HU, ID, IN, MX, MY, PE, PH, PL, TH, TR and ZA; euro area = BE, DE, ES, FR, IT and NL; other AEs = AU, CA, CH, JP and SE. Data for China until December 2025.

Graph 2.A: Magnificent 7 =Alphabet, Amazon, Apple, Meta, Microsoft, Nvidia and Tesla. For the Magnificent 7 and the rest of the S&P 500, market capitalisation-weighted average. The composition of the S&P 500 is fixed on 5 March 2026.

Graph 2.B: For each sector, average of the respective index constituents' returns.

Graph 2.C: Measured by the S&P 500 Growth and Value indices, respectively.

Graph 3.A: Dispersion index = Cboe S&P 500 Dispersion Index; oil VIX = Cboe Crude Oil ETF USO VIX Index; VIX = Cboe Volatility Index.

Graph 3.B: Five-day moving average.

Graph 3.C: EU natgas = Title Transfer Facility (TTF); precious metals index = simple average across the spot prices of palladium, platinum and silver.

Graph 4.A: Spreads of ICE BofAML index yields to OIS rates with matched maturities. Ten-year OIS for US investment grade, five-year for the rest. The USD OIS rate is based on SOFR starting on 11 December 2018. The EUR OIS rate is based on ESTR starting on 11 October 2019. For US investment grade, the box plot shows data between 31 July 2008 and 5 March 2026, due to data availability of USD 10-year OIS.

Graph 4.B: Twelve-month moving sum.

Graphs 6.A and 6.B: 1 December 2025 = beginning of the review period; 27 February 2026 = last observation prior to the start of the conflict in the Middle East; 5 March 2026 = end of the review period.

Graph 6.C: Ten-year minus two-year government bond yield.

Graph 7.A: Simple average. Asia = CN, ID, IN, MY and TH; EMEA = HU, PL and ZA; Latin America = BR, CL, CO, MX and PE. Five-day moving average.

Graph 8.C: Asia = ID, IN, LK, MN, MY, PH, PK, TH, TW and VN; Latin America = BR, CL, CO, and MX; other EMEs = GH, HU, KE, LB, MK, PL, QA, RO, RS, RU, SA, TR, UA and ZA.

Sarah Bell sarah.bell@bis.org
Matthieu Chavaz matthieu.chavaz@bis.org
Boris Hofmann boris.hofmann@bis.org
Daniel Rees daniel.rees@bis.org
Matthias Rottner matthias.rottner@bis.org

Evolving approaches to monetary policy communication in the face of uncertainty: fan charts, scenarios and guidance[1]

The elevated and rapidly shifting nature of uncertainty in recent years has posed challenges for central bank communication. In particular, central banks have had to reconsider how to communicate the economic outlook and the monetary policy response. They have expanded their communications toolkit to explain uncertainty about the outlook and are increasingly using scenario analysis to complement traditional tools like fan charts and qualitative risk discussions. Fewer central banks are providing descriptive policy guidance; instead more are publishing their own policy rate projections, often in the context of alternative scenarios.

JEL classification: D81, E17, E58

Communication has played a more central role in monetary policy in recent decades. Explicit price stability objectives, often in the form of inflation targets, combined with central bank independence, have been accompanied by greater expectations of transparency and accountability. Communicating the central bank's views about the economic outlook and the reasoning behind its policy decisions helps to address these expectations. At the same time, communication can enhance monetary policy effectiveness. It can help anchor inflation expectations at the inflation target, supporting economic stability more broadly. And it can help align interest rate expectations with the central bank's envisaged policy rate path, reinforcing monetary policy transmission.[2]

Monetary policy communication has become more challenging in recent years in the wake of unprecedented shocks (eg the Covid-19 pandemic) and increasing geopolitical tensions. As a result, central banks have been placing greater emphasis on uncertainty in their public communications. While in 2006 most monetary policy statements did not mention uncertainty, in 2025 the share of words related to uncertainty in these statements rose to up to about 15% (Graph 1.A). Similarly, the share of speeches by central bank officials that mentioned uncertainty as a key theme rose to around 40% (Graph 1.B). At the same time, a growing body of literature has

[1] The views expressed in this publication are those of the authors and not necessarily those of the BIS or its member central banks. We thank Anamaria Illes for excellent research support. For insightful comments, we thank Eduardo Amaral, Claudio Borio, Gaston Gelos, Emmanuel Kohlscheen, Marco Lombardi, Benoît Mojon, Alejandrina Salcedo, Andreas Schrimpf, Frank Smets, Ilhyock Shim and Hyun Song Shin. All remaining errors are our own.

[2] See Blinder et al (2008) for a more detailed discussion.

Key takeaways

- *The elevated and rapidly shifting nature of uncertainty in recent years has affected the way central banks communicate uncertainty around the economic outlook and the monetary policy response.*
- *More central banks are using scenarios to illustrate the implications of specific risks, in addition to traditional tools like fan charts and qualitative risk discussions.*
- *Central banks are providing less descriptive policy guidance, and more are publishing their own policy rate projections, often in the context of alternative scenarios.*

examined the implications of uncertainty for the communication strategies of individual central banks.[3]

In particular, central bank communication about the future is becoming more complex in a context of heightened uncertainty, especially where the implications of possible developments are hard to quantify. Definitive statements about the economy's evolution will often be wrong. Yet overemphasising uncertainty can obscure the public's understanding of policymakers' views about the economic outlook. High uncertainty can contribute to large forecast errors and increases the likelihood that policy guidance will need to be revised as conditions change, with costs to the central bank's reputation and credibility. Conversely, high uncertainty may increase the value of clarifying the central bank's reaction function, thereby reducing uncertainty about how the central bank will respond to economic developments.

Against this background, this article explores how central banks address uncertainty in their public communication. Specifically, we examine the tools central banks use to discuss uncertainty around the economic outlook and those they use to convey the policy outlook, as well as how their use of these tools has changed over time. To this end, we document the evolution of central bank communication approaches over the past 20 years based on new hand-collected data from 25 central banks operating under inflation targeting or similar policy regimes.[4] Drawing on published monetary policy statements and monetary policy reports, we identify the main tools used to address uncertainty around the economic outlook and influence uncertainty about the policy outlook.[5] We then document how the use of these tools has evolved across central banks based on the frequency of use in four specific years – 2006, 2012, 2019 and 2025. In addition, we use a standard New Keynesian model to illustrate some of the challenges that different communication tools pose under uncertainty.

The article first lays out a conceptual framework, mapping different types of uncertainty into monetary policy communication tools. It then assesses changes in

[3] See eg Bernanke (2024, 2025), Hunter (2024), Lombardelli (2025), Hernández de Cos (2025), Bauer et al (2025), Amaral et al (2025) and Lane (2024).

[4] We include the central banks of Australia, Brazil, Canada, Chile, Colombia, Czechia, euro area, Hungary, Iceland, India, Israel, Japan, Korea, Mexico, New Zealand, Norway, Peru, Poland, Romania, South Africa, Sweden, Switzerland, Thailand, the United Kingdom and the United States.

[5] Central banks also use other important communication vehicles to convey their views and the rationale behind their decisions, notably press conferences, publication of minutes of monetary policy meetings, and speeches. However, the monetary policy report and monetary policy statements are the primary vehicles for communicating the central bank's views on uncertainty about the economic outlook and the outlook for monetary policy.

Uncertainty in monetary policy statements and speeches[1]

In per cent

Graph 1

A. Share of words related to uncertainty in monetary policy statement[2]

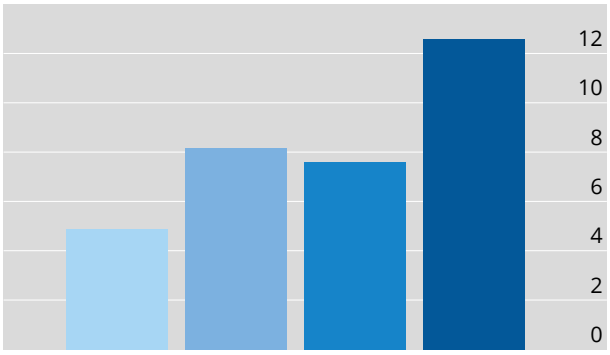

B. Central bank speeches highlighting uncertainty[3]

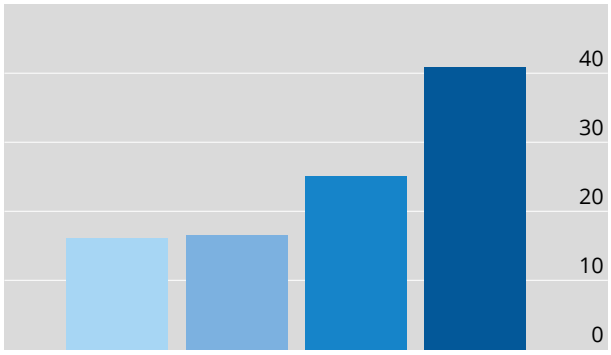

[1] See endnotes for details. [2] Simple average across central banks of the number of words related to risks or uncertainties about the outlook as a share of total number of words in the monetary policy statement. Only central banks operating an inflation targeting or closely related policy regime during the relevant time frame are included. [3] Share of central bank speeches that include "uncertainty" as a key message, relative to all speeches.

Sources: Park et al (2025); central bank reports; BIS; authors' calculations.

central banks' approaches to address uncertainty about the economic outlook, as well as changes in the communication of the monetary policy response. We end with a brief discussion of the implications of our analysis for the future of central bank communication.

Communicating under uncertainty: a simple taxonomy

Several taxonomies of uncertainty and risks have been proposed in the academic literature, but none relate specifically to communication. Bernanke (2007) proposes three main categories of uncertainty referring to the state of the economy, the structure of the economy and private sector expectations. Knight (1921) distinguishes between measurable uncertainty (risk) and unmeasurable (true) uncertainty, with the latter referred to in the literature as Knightian uncertainty.[6]

For assessing the implications of uncertainty for central bank communication tools, an alternative taxonomy focusing on its form is useful. In this article, we distinguish between two forms of uncertainty. The first is *general*, arising from the inherent unpredictability of the future, as the economy rarely evolves in line with the baseline forecast.[7] While general uncertainty is always present, at different times it may be perceived as higher or lower. The second form of uncertainty is *specific*,

[6] See eg Kozicki and Vardy (2017) for a more detailed discussion.

[7] These forecast errors can occur because the central bank (and the private sector) has incomplete information about the model of the economy or because of unanticipated developments.

stemming from identifiable risks or developments that can affect the economy, and hence the monetary policy stance, such as the introduction of trade tariffs.

Central bank communication addresses these two uncertainties in distinct ways. Central banks typically seek to *convey* the uncertainty around the economic outlook by illustrating the ways in which their forecasts could be wrong and explaining the implications of possible deviations. In contrast, central bank communication often aims to *reduce* policy response uncertainty by describing how monetary policy is likely to react to evolving economic conditions.

Central banks have deployed a range of communication tools to address these different aspects of uncertainty. The tools have distinct benefits and drawbacks, because each is generally best suited to address certain forms of uncertainty (Table 1).

Methods to address economic outlook and policy response uncertainty Table 1

Uncertainty about	General uncertainty	Specific uncertainty
Economic outlook	Fan charts	Alternative scenarios, risk discussions
Policy response	Descriptive guidance about policy path, policy rate projections	Scenarios with policy rate response

Source: Authors' elaboration.

To describe uncertainty about the economic outlook, central banks have used three main tools:

- Fan charts provide a graphical representation of the general uncertainty around the forecast for inflation (and sometimes GDP growth and other variables). They depict confidence intervals for certain deviations from the central forecast, typically estimated from past forecast errors.
- Alternative scenarios provide quantitative assessments of the macroeconomic implications of specific sources of risk. These can include, for instance, the economic impact of hypothetical exogenous shocks, alternative paths of key conditioning variables or structural change in the economy.
- Risk discussions qualitatively highlight specific sources of uncertainty. The amount of detail they contain can vary from a brief list of key risks to an extensive discussion of the triggers for each specific risk and their expected economic impact.[8]

Some central banks provide information about future policy to reduce uncertainty about their expected future decisions and improve understanding of their reaction function. Here too three tools are prevalent:

- Descriptive guidance provides an explanation of the expected monetary policy stance, influencing the degree of general uncertainty among the public about policy and the reaction function. This guidance can indicate the central bank's assessment of the most likely future policy stance (Delphic

[8] For example, the "Key risks to the outlook" section of the latest Reserve Bank of Australia monetary policy report highlights issues such as the strength of household demand, labour market tightness and global growth prospects. The Brazilian central bank's report distinguishes between upside and downside inflation risks (eg due to a prolonged de-anchoring of inflation expectations or a steeper global slowdown).

guidance) or involve a commitment to adhere to a specific path, either for a set period or as long as certain conditions are met (Odyssean guidance).[9]

- Policy rate projections provide a quantitative indication of the expected policy path, typically in the form of a chart in the central bank's monetary policy report.[10] These projections can also influence the extent of general policy uncertainty, although relative to descriptive guidance they are more tightly linked to the central forecast, which in turn provides greater clarity about the central bank's reaction function.
- Scenarios with a policy response illustrate the central bank's expected reaction in the specific context of the scenario. Such scenarios hence influence uncertainty about the policy response in specific situations and can, more broadly, further clarify the reaction function.

The mapping in Table 1 suggests that central banks use tools to address both general and specific forms of uncertainty in communication about the economic outlook and policy response. The table also shows how the specific mix of tools used may be adjusted based on the form of perceived uncertainty confronting central banks.

Communicating uncertainty around the economic outlook

The tools used to communicate uncertainty around the economic outlook vary across central banks (Graph 2). Fan charts are the most common tool, published by almost 75% of the 25 central banks in our sample in 2025. Qualitative risk discussions accompanying the publication of the forecast are also a popular tool to convey risks around the outlook, used by almost 70% of the assessed institutions. Fewer central banks, around 45%, employ scenario analysis to quantitatively illustrate risks around the baseline forecast.[11]

The tools used have also evolved over time (Graph 2). The most noteworthy change concerns scenario analysis. Published by only about 25% of assessed central banks before the Covid-19 pandemic, its use has risen substantially since then. Qualitative risk discussions, already presented by about half of the assessed central banks in 2006, became even more common after the Great Financial Crisis. Their use has remained broadly stable since. Fan charts, pioneered in 1993 by the Bank of England in its inaugural *Inflation Report*, saw steady adoption from the start of our sample and were presented by around 75% of assessed central banks in 2019. Unlike other recent surveys of central bank communication approaches (eg Bauer et al (2025)), ours – which includes a larger sample of central banks – indicates that the

[9] See Campbell et al (2012).

[10] We consider here only projections of the policy rate (or closely related short-term interest rate) made by the central bank itself and not technical assumptions (eg a constant interest rate assumption) or projections based on an external source (eg financial market pricing).

[11] Some central banks present scenarios irregularly. For example, in recent years the Reserve Bank of Australia has presented them only in its May and November monetary policy reports, while the Bank of Canada presented scenarios in April 2025 to illustrate the effects of trade tariffs.

Central bank approaches to communicating economic outlook uncertainty[1]

As a percentage of central banks

Graph 2

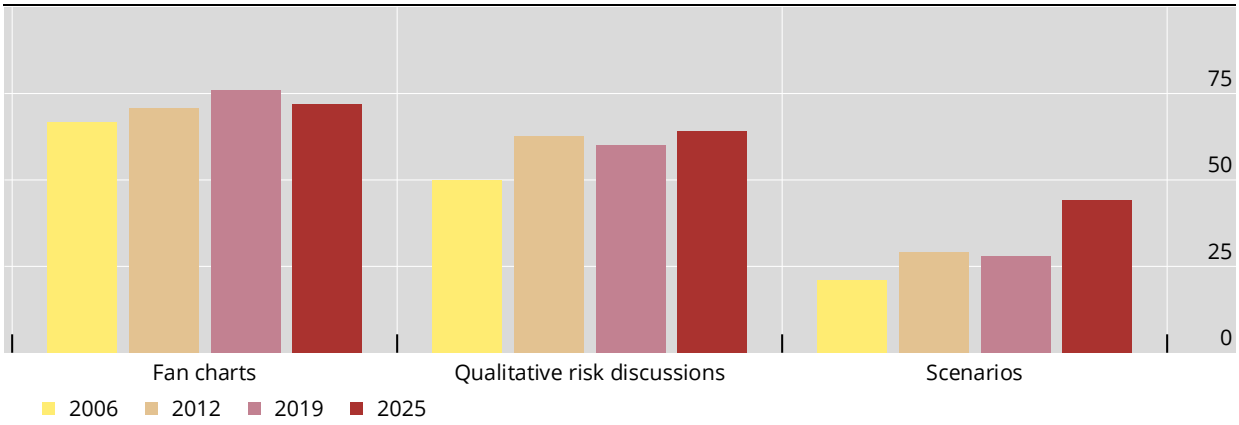

[1] Data obtained from monetary policy reports of central banks or similar publications. Only central banks operating an inflation targeting or closely related policy regime during the relevant time frame are included. See endnotes for details.

Sources: Central banks; authors' calculations.

use of fan charts has dipped only marginally in recent years. However, their prominence within central bank reports has diminished in some cases.[12]

The greater prevalence of scenario analysis and relative de-emphasis of fan charts by some central banks between 2019 and 2025 partly reflect the strengths and weaknesses of the two tools in addressing the perceived uncertainties prevailing since the outbreak of the Covid-19 pandemic.

Fan charts highlight the general uncertainty in central banks' economic forecasts. They provide a concise way to communicate the central bank's assessment of the likelihood of key economic outcomes, such as overshooting the inflation target or entering a recession, without the need to detail every contributing factor. They can also help to manage expectations about the central bank's knowledge of the future.

However, if too wide, fan charts risk diluting the central bank's messages about the outlook and raising questions about its forecasting ability. This problem can become acute if large shocks give rise to large forecast errors, as past errors typically form the basis for determining the width of the fan chart. Graph 3 illustrates this with fan charts constructed based on a standard New Keynesian model estimated on aggregate G7 country data (see the online annex for details). The graph shows how the fan chart for GDP growth inflates when the Covid-19 pandemic observations are included (Graph 3.A) compared with when the model features an adjustment for these large shocks (Graph 3.B).

Another drawback with fan charts is their inability to highlight specific risks facing the economy or to capture bi- or multi-modal distributions of possible future outcomes. This limitation has become particularly relevant in recent years due to the presence of large perceived uncertainties in the form of distinct, plausible paths for

[12] The Central Bank of Norway, for instance, presents its main economic forecasts at the beginning of its monetary policy report as point forecasts and shows fan charts towards the end of the report as part of a discussion of historical forecast errors. Until recently it had suspended fan charts in the wake of the Covid-19 pandemic.

Large shocks can widen fan charts

In per cent

Graph 3

A. GDP growth fan charts that include the Covid-19 pandemic are much wider…[1]

B. …than they are with an adjustment for Covid-19[2]

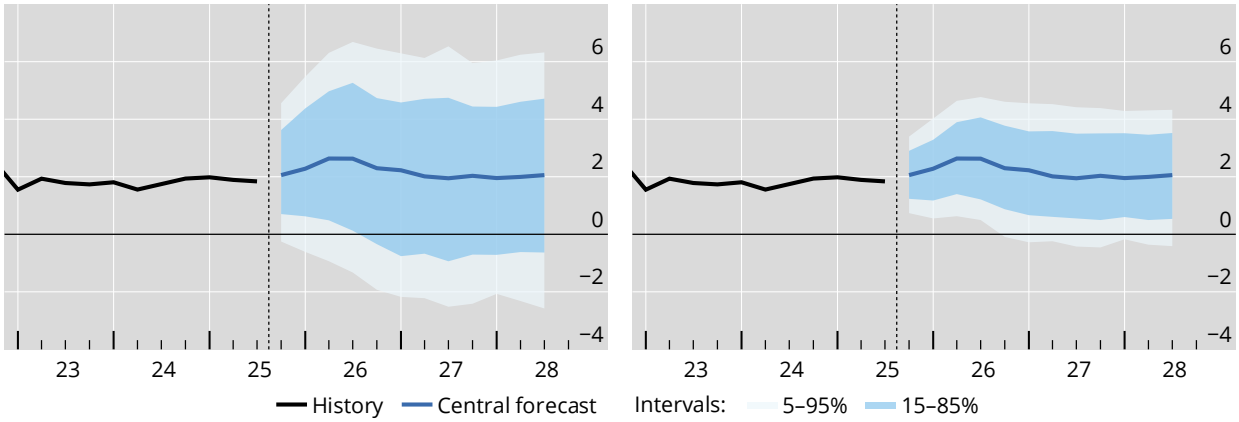

[1] Fan chart are constructed based on the estimation of the New Keynesian model described in the online annex without any adjustment for Covid-19. [2] Forecast of real GDP growth and fan charts are constructed from the same model featuring a Covid-19 adjustment in the estimation.

Sources: National data; authors' calculations.

the economy. A prominent example occurred in 2020, when uncertainty arose about whether there would be a timely vaccine for Covid-19. Another was the inflation surge of 2021–22, when there was the question about whether the inflation shock would prove transitory, persistent or potentially even trigger a wage-price spiral. In these cases, a fan chart centred on a baseline forecast that averaged out the alternatives may provide a poor guide to the uncertainty around the economic outlook. Instead, central banks often find it useful to rely on tools that can draw attention to specific risks and describe their economic consequences.

Scenario analysis is one such tool. It is particularly useful for crafting narratives around complex economic dynamics to help clarify how specific risks may play out. A useful example is to show the trade-offs associated with responding to higher inflation and how these depend on the underlying causes. A central bank could illustrate this by presenting two scenarios, one in which households save less of their income than expected (a positive "aggregate demand shock") and one where productivity grows more slowly than usual (a negative "aggregate supply shock"). Although both scenarios raise inflation (Graph 4.A), the path of GDP differs sharply (Graph 4.B). This exercise presents a concise visual illustration of how a given inflation path can have different growth and policy implications.

Central banks also commonly use scenario analysis to explain how structural shifts can affect the economic outlook. A typical example would be to compare the standard effects of a positive aggregate demand shock with an alternative where higher inflation leads firms to adjust their prices more often. More frequent price changes can strengthen the relationship between economic slack and inflation. Such a scenario can be constructed by applying the same sequence of aggregate demand shocks as in Graph 4.A to a model with a steeper Phillips curve (Graph 4.C).

The examples in Graph 4 also illustrate two key challenges central banks face when presenting scenarios: selecting and scaling them. When the economy faces

large, discrete and identifiable sources of uncertainty, the choice of scenarios can be straightforward. For example, in early 2025 many central banks presented scenarios illustrating the effects of trade tariffs. At other times, the central bank must decide which of the many sources of uncertainty to highlight. Because central banks rarely present more than one or two scenarios at a given time, this choice can raise the profile of particular risks, possibly at the expense of other, equally impactful ones. Indeed, this may explain why some central banks publish scenarios on an ad hoc basis rather than releasing them routinely.

The scaling – or calibrated severity – of a given scenario is typically at the discretion of the central bank. In making this choice, central banks must balance assessments about the likelihood of the scenario with the information it provides. Scenarios that are scaled so that they cluster around the central forecast are arguably more likely to occur but may provide relatively little information value. Still, these scenarios can be useful for central banks, as they show the most likely ways in which economic outcomes – and hence policy settings – could differ from the central forecasts.

Scenarios scaled to present more meaningful departures from the central forecasts may allow the central bank to provide more information about material risks to the outlook. The drawback is that such scenarios may seem too unlikely to matter in practice. The lower aggregate supply scenario in Graph 4.A requires an unusually large sequence of shocks, while the increase in the steepness of the Phillips curve in Graph 4.C far exceeds typical empirical estimates of this parameter. Calibrating shock

Scenarios can help central banks explain complex economic concepts[1]

In per cent

Graph 4

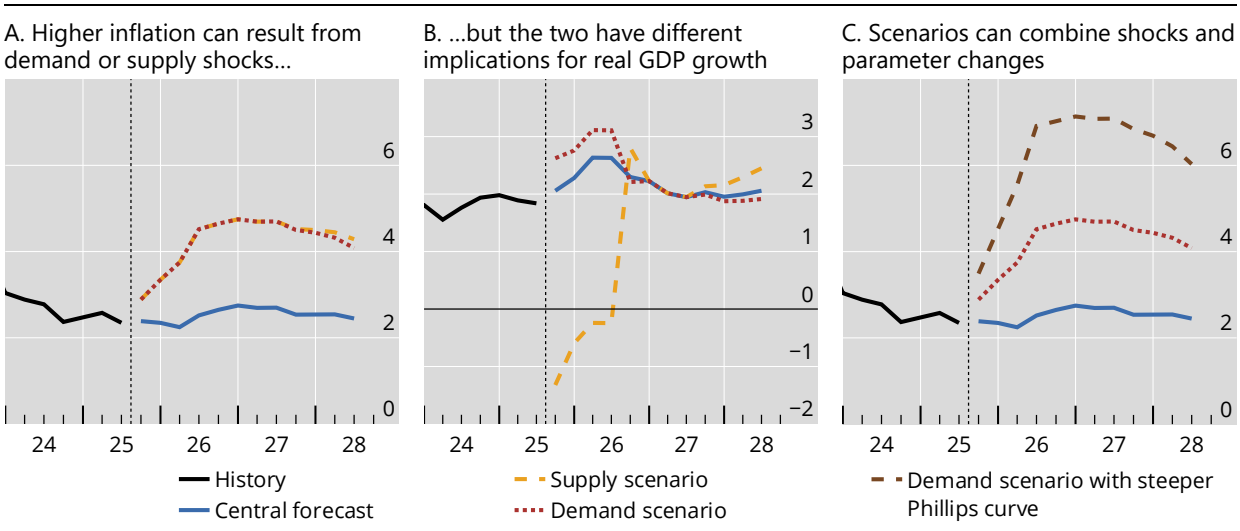

[1] Forecasts and scenarios are constructed using the New Keynesian DSGE model described in the online annex. Scenarios are constructed by applying the model's total factor productivity shocks (supply scenario) and demand shocks (demand scenario). The same sequence of demand shocks is applied to both the demand scenario and the demand scenario with steeper Phillips curve.

Sources: National data; authors' calculations.

sequences or changes in Phillips curve parameters to historical estimates would yield much smaller departures from the central forecast.[13]

Communicating the policy response

In 2025, providing explicit guidance about future policy rates in any form was relatively uncommon (Graph 5, red bars). Around half of the central banks in our sample did not provide descriptive guidance or publish a policy rate path. And no individual form of guidance was used by more than a third of them.

Communication practices about the policy response have shifted notably over time. Until recently, many central banks provided descriptive policy rate guidance.[14] This was particularly widespread in the 2010s, when more than 55% of central banks in our sample provided guidance, compared with only 20% in 2025 (Graph 5). Meanwhile, the share of central banks publishing a forecast path for their own policy rate or presenting scenarios with an accompanying policy rate path has risen steadily. That said, these tools are still employed by only a minority of our sample of central banks.

Changes in communication practices about the policy outlook also reflect, in part, perceived shifts in the nature of uncertainty. The widespread use of descriptive guidance in the 2010s was closely tied to the post-GFC low interest rate environment, when such communication helped to lower the public's interest rate expectations and

Communicating the policy outlook[1]

As a percentage of central banks — Graph 5

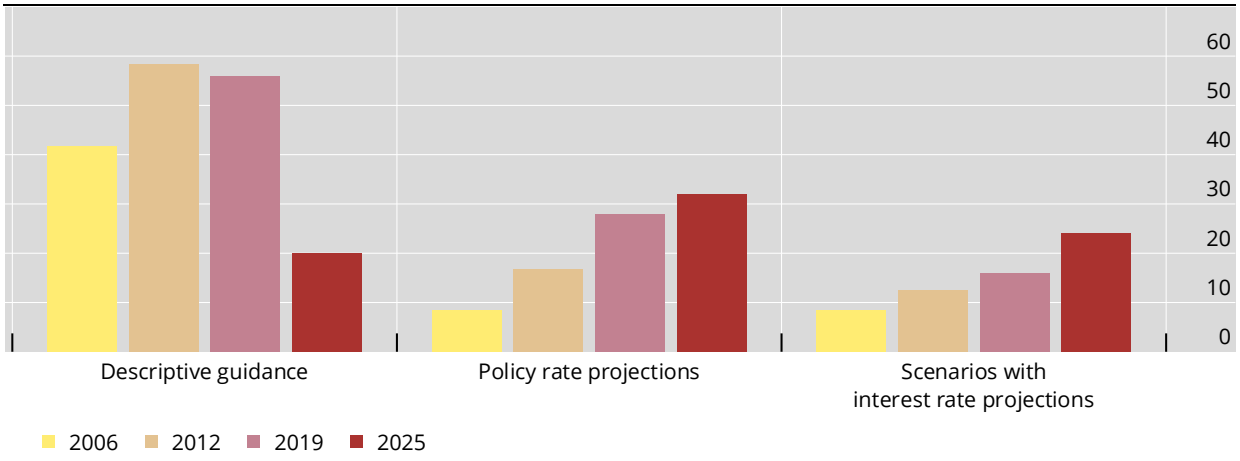

[1] Descriptive guidance in the monetary policy statement of central banks. Policy rate projections and scenarios with interest rate projections from monetary policy reports of central banks or similar publications. Only central banks operating an inflation targeting or closely related regime during the relevant time frame are included. See endnotes for details.

Sources: Central banks; authors' calculations.

[13] A further challenge with more extreme scenarios is that they can be hard to implement with standard analytical tools and models and may call for the development of more advanced techniques. See eg Kase et al (2025) and Koyuncu et al (2026) for recent advances in structural and time series modelling.

[14] Defined as a qualitative indication of the likely direction for future policy in the monetary policy statement.

hence longer-term yields. At that time, perceived risks were skewed to the downside. Developments calling for rapid tightening were seen as unlikely, while the effective lower bound on policy rates raised concerns that central banks could respond less forcefully to negative inflation surprises than to positive ones. In this environment, the perceived costs – chiefly reduced policy flexibility – of lowering general uncertainty about future policy by signalling that interest rates would remain low (potentially subject to certain conditions) were judged small relative to the benefits of additional stimulus.[15]

The post-Covid inflation surge revealed the drawbacks of descriptive guidance. Even when such guidance is explicitly tied to specific economic developments, the public often views it as an unconditional commitment to a certain policy path. This may constrain the central bank's perceived room for manoeuvre when economic conditions suddenly change, or it may erode the central bank's credibility if the public feels it has reneged on its policy commitments.[16]

In place of descriptive guidance, many central banks have recently emphasised the "data dependent" nature of their decisions. In doing so, they have underscored that the general uncertainty about the economic outlook implies a degree of uncertainty about future policy. The more widespread publication of the central bank's policy rate projections as well as variations of the path under different scenarios is consistent with this greater emphasis on the conditionality, or data dependence, of the policy rate outlook.

Publishing the central bank's own policy rate projection also provides a degree of quantitative forward guidance, which may help to shape financial market expectations (Rudebusch and Williams (2008)). Unlike some forms of descriptive guidance, however, policy rate projections are explicitly tied to the central forecast, underscoring their conditionality. Indeed, one possible use of policy rate projections is to present an empirically testable link between the policy rate forecast and the forecast for inflation and real economic activity (Hofmann and Xia (2022)). In this way, interest rate paths may help reduce uncertainty about the central bank's reaction function and its assessment of monetary policy transmission.[17]

However, there are also potential drawbacks associated with publishing the central bank's own policy rate projection. Financial market participants may focus on it excessively, reducing the information content of market prices and making it hard for the central bank to deviate. Conversely, the central bank may face reputational risks if deviations from the projected interest rate path create the perception that it is unable to forecast its own policy rate accurately. This may explain why, despite an uptrend in the use of the tool, the number of central banks that publish a policy rate projection is still relatively small.[18]

Scenarios with variations in the policy rate path may also help to highlight the conditionality of the baseline path and further clarify the central bank's reaction

[15] See Committee on the Global Financial System (2019).

[16] Filardo and Hofmann (2014) provide an early assessment of different forms of forward guidance at the zero lower bound. See also Reserve Bank of Australia (2022) for a detailed discussion of the RBA's experience with forward guidance.

[17] Inferring the central bank's reaction function from projections of the policy rate and relevant macro variables may still be complicated if the public is unsure of the central bank's estimate of the long-run natural interest rate.

[18] For a more detailed discussion of these issues, see Hofmann and Xia (2022) and the references therein.

function. By linking the policy response to a particular scenario, this approach can reinforce the message that future policy depends on economic developments.[19] Such scenarios can further help markets and the public to anticipate how potential developments would be factored into policy decisions. It can also underscore the central bank's willingness to take necessary measures to achieve its objectives, helping to anchor inflation expectations.

However, this communication tool also poses challenges. A fundamental one is that it may be hard for the central bank to predict how policy would evolve in a hypothetical situation. A common approach is to apply a "policy rule" that maps the policy stance to macroeconomic variables, such as inflation or GDP growth. Often, these rules are estimated based on the central bank's past behaviour. While simple and transparent, it can be difficult to capture succinctly the full range of considerations that policymakers take into account, such as financial conditions. Hence, estimated policy rules may provide a poor guide to the central bank's actual actions if faced with a given scenario.

A second issue with presenting scenarios with a policy response is that policy rates both affect other economic variables and are affected by them. This circular relationship means that the choice of policy rule can fundamentally shape the nature of the scenario that is communicated. For example, Graph 6 shows a scenario featuring a demand shock that, holding other factors equal, would raise inflation by 2 percentage points across the next two years. Higher inflation, however, naturally prompts a higher policy rate, which reduces the rise in inflation by half. Instead of representing a surge in inflation and GDP growth that poses a material risk to the forecast (Graph 6, orange lines), the scenario displays a relatively modest forecast

Communicating the policy path poses challenges in scenario analysis[1]

In per cent

Graph 6

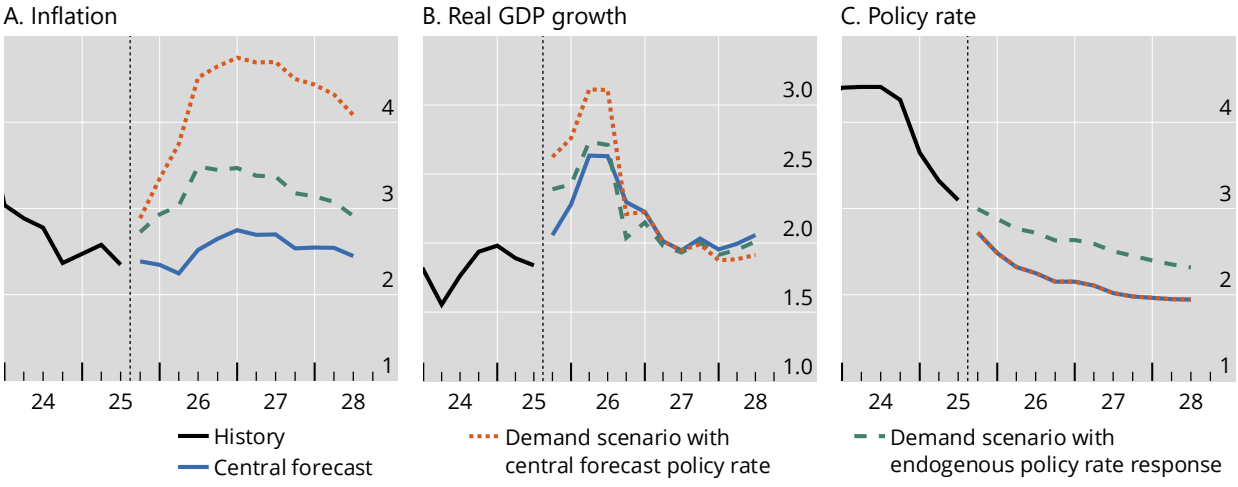

[1] Forecasts and scenarios are constructed using the New Keynesian DSGE model described in the online annex. Scenarios are constructed by applying the model's demand shocks (demand scenario). The same sequence of demand shocks is applied to both the scenario with central forecast policy rate and the scenario with endogenous policy rate response.

Sources: National data; authors' calculations.

[19] Such information may be particularly useful for extreme scenarios that may prompt the central bank to depart from its usual policy reaction function to limit the likelihood of particularly adverse outcomes.

miss (green lines). While presenting the scenario in this way can help the central bank reinforce its commitment to do what is required to meet its inflation objective, showing an inflation path that is close to the central scenario might lead some observers to question why tighter policy is necessary.

These challenges may explain why many central banks that present scenarios assume that policy follows the same path as the central forecast. While scenarios that do not present an endogenous policy rate path cannot provide concrete guidance about the central bank's reaction function, they can highlight the costs of not adjusting policy as conditions change, creating a rationale for action if the scenario actually arises. That said, there is a risk of suggesting that the central bank would be unable or unwilling to meet its inflation objectives if the scenario occurred, potentially weakening inflation expectations. To mitigate this, the central bank could present the scenario with both a policy reaction function and a rate path that matches the central forecast.

Conclusion

Central banks' approaches to addressing uncertainty about the economic outlook and the policy response in their public communication have evolved over time. This article documents this evolution and links it to the characteristics of different communication tools and the changing nature of uncertainty. A key conclusion is that no single tool is right for all situations. Each has distinct strengths and weaknesses that vary with the context, and none can address all forms of uncertainty. Nonetheless, understanding these characteristics can help central banks use the most appropriate tool, or combination of tools, for communicating the perceived uncertainty they face at a given time.

Looking ahead, central banks will likely need to continue adapting their communication approaches to address uncertainty as it evolves. In doing so, three tasks will be central: (i) identifying clear objectives and measures of successful communication; (ii) assessing the effectiveness of existing approaches, and potentially innovative new tools, against these objectives; and (iii) deploying communication approaches to support desired policy outcomes.

Endnotes

Graph 1.A: Sample includes AU, BR, CA, CH, CL, CO, CZ, EA, GB, HU, IL, IN, IS, JP, KR, MX, NO, NZ, PE, PL, RO, SE, TH, US and ZA.

Graph 1.B: Sample includes AU, CA, CH, DE, EA, ES, FR, GB, HK, IE, IN, IT, JP, MY, NO, PH, SE, SG, US and ZA.

Graph 2: Sample includes AU, BR, CA, CH, CL, CO, CZ, EA, GB, HU, IL, IN, IS, JP, KR, MX, NO, NZ, PE, PL, RO, SE, TH, US and ZA.

Graph 5: Sample includes AU, BR, CA, CH, CL, CO, CZ, EA, GB, HU, IL, IN, IS, JP, KR, MX, NO, NZ, PE, PL, RO, SE, TH, US and ZA.

References

Amaral, E, F Avalos and I Shim (eds) (2025): "Monetary policy decision-making and communication under uncertainty", *BIS Papers,* no 163, December.

Bauer, M, T Berge, G Fiori, F Loria and M Zhong (2025): "Accounting for uncertainty and risks in monetary policy", *Finance and Economics Discussion Series,* no 2025-73.

Bernanke, B (2007): "Monetary policy under uncertainty," speech delivered at the 32nd Annual Economic Policy Conference, Federal Reserve Bank of St. Louis, 19 October.

——— (2024): *Forecasting for monetary policy making and communication at the Bank of England: a review*, Independent Evaluation Office, Bank of England, April.

——— (2025): *Improving Fed communications: a proposal from Ben Bernanke*, Brookings, May.

Blinder, A, M Ehrmann, M Fratzscher, J De Haan and D-J Jansen (2008): "Central bank communication and monetary policy: a survey of theory and evidence", *Journal of Economic Literature*, vol 46, no4.

Blinder, A, C Goodhart, P Hildebrand, D Lipton and C Wyplosz (2001): "How do central banks talk?", *Geneva Reports on the World Economy*, no 3, Centre for Economic Policy Research.

Campbell, J, C Evans, J Fisher and A Justiniano (2012): "Macroeconomic effects of Federal Reserve forward guidance", *Brookings Papers on Economic Activity*, Spring.

Committee on the Global Financial System (2019): "Unconventional monetary policy tools: a cross-country analysis", *CGFS Papers,* no 63, October.

Filardo, A and B Hofmann (2014): "Forward guidance at the zero lower bound", *BIS Quarterly Review*, March.

Hernández de Cos, P (2025): "Delivering on central bank mandates in a changing world", speech at the Bank of Mexico 100th Anniversary Conference, Mexico City, 26 August.

Hofmann, B and D Xia (2022): "Quantitative forward guidance through interest rate projections", *BIS Working Papers*, no 1009, March.

Hunter, S (2024): "Shedding light on uncertainty: using scenarios in forecasting and policy", speech at the University of Adelaide South Australian Centre for Economics Studies Lunch, Adelaide, 13 December.

Kase, H, L Melosi and M Rottner (2025): "Estimating nonlinear heterogeneous agent models with neural networks", *BIS Working Papers*, no 1241, January.

Knight, F (1921): *Risk, uncertainty and profit,* Hart, Schaffner & Marx; Houghton Mifflin Co.

Koyuncu, B, B Kwon, M Lombardi, F Perez-Cruz, and H Shin (2026): "BISTRO: a general-purpose oracle for macroeconomic time series", *BIS Quarterly Review*, March.

Kozicki, S and J Vardy (2017): "Communicating uncertainty in monetary policy", *Bank of Canada Staff Discussion Papers*, no 2017-14, November.

Lane, P (2024): "Monetary policy under uncertainty", keynote speech at the Bank of England Watchers' Conference, London, 25 November.

Lombardelli, C (2025): "What if things are different?" speech at the Bank of England Watchers' Conference, London, 25 November.

Park, T, F Perez-Cruz and H Shin (2025): "Mapping the space of central bankers' ideas", *BIS Working Papers*, no 1299, October.

Reserve Bank of Australia (2022): *Review of the RBA's approach to forward guidance*, https://www.rba.gov.au/monetary-policy/reviews/approach-to-forward-guidance/index.html.

Rudebusch, G and J Williams (2008): "Revealing the secrets of the temple: the value of publishing central bank interest rate projections", in Campbell (ed), *Asset prices and monetary policy*, National Bureau of Economic Research.

Prashant R Babu
prashant.babu@bankofengland.co.uk

Michael Chui
michael.chui@bis.org

Costas Stephanou
costas.stephanou@bis.org

The rise and risks of synthetic risk transfers[1]

Synthetic risk transfers (SRTs) are used by banks for capital and credit risk management. They have grown significantly in recent years but remain small relative to bank balance sheets. European banks have traditionally dominated SRT issuance, but their North American peers have seen an increase in issuance more recently. SRT-related risks appear to be modest at present. However, this may change as the market expands, structures become more complex and banks rely more heavily on non-bank financial institutions for credit protection. Limited and fragmented information heightens the potential for SRT-related risks to build undetected, highlighting the need for enhanced monitoring of risks for individual banks and from a system-wide perspective.

JEL classification: G18, G21, G22, G23, G28.

Synthetic risk transfers (SRTs), once a niche securitisation tool for capital and credit risk management, have grown rapidly in recent years. These instruments allow banks to retain ownership of a portfolio of loans while transferring part of the credit risk to institutional investors by means of derivatives or guarantees. The growth in SRT issuance has facilitated the redistribution of risk from banks to non-bank financial institutions (NBFIs), thereby deepening the linkages between the two sectors.

This article provides an overview of SRT trends and their key drivers and assesses the potential risks of these instruments. A key contribution is to examine whether and under what conditions SRTs may have system-wide risk implications by creating new – or amplifying existing – channels of contagion in stress. The analysis builds on other recent SRT work (BCBS (2026); Cortes et al (2025); ESRB (2025)) and informs the debate on evolving business models and interlinkages between banks and NBFIs.

Three findings emerge from the analysis. First, SRTs have grown significantly in recent years. Issuance has increased fivefold since 2016, providing protection to loan portfolios of almost €800 billion as of end-2024. Around eight new banks have issued SRTs every year since 2016, raising the cumulative number of issuers above 100. This expansion, which has allowed banks to support lending to the real economy by transferring credit risk to NBFIs, was driven by regulatory and supervisory changes, bank balance sheet optimisation and increased demand from institutional investors.

[1] The views expressed in this publication are those of the authors and not necessarily those of the BIS or its member central banks. We are grateful to the International Association of Credit Portfolio Managers (IACPM) for organising industry outreach and to Burcu Erik for excellent research assistance. We also thank Puriya Abbassi, Jeremy Brizzi, Sebastian Doerr, Mathias Drehmann, Gaston Gelos, Maciej Grodzicki, Tamar Joulia-Paris, Ulf Lewrick, Benoît Mojon, Richard Murdoch, Woojung Park, Jitendra Patil, Daniel Rees, Bertrand Rime, Brendan Rowan, Andreas Schrimpf, Hyun Song Shin, Nikola Tarashev and Chen-Helen Zhang for helpful comments and suggestions.

Key takeaways

- *The use of SRTs for bank capital and credit risk management has grown significantly in recent years, though the market remains concentrated and small relative to bank balance sheets.*
- *SRT-related risks appear to be modest at present. However, this may change as the market expands, structures become more complex and banks rely more heavily on NBFIs for credit protection.*
- *Limited and fragmented information heightens the potential for SRT-related risks to build undetected, highlighting the need for enhanced monitoring of risks for banks and from a system-wide perspective.*

Notwithstanding this growth, the SRT market remains small and concentrated. These instruments protected around 2% or less of total bank loans in the European Union (EU), United States, United Kingdom and Canada at end-2024. They provided capital relief of around 43 basis points of Common Equity Tier 1 (CET1) for SRT-issuing banks, which is modest compared with sector-wide average CET1 levels of 14–16% in these jurisdictions. The SRT market is concentrated in terms of issuers (mostly EU banks), asset type (mostly corporate loans) and investors (mostly credit funds and asset managers), although it has become more diverse in recent years.

Second, SRT-related risks appear modest at present but could increase as the market expands. SRTs demonstrated resilience during recent episodes of market turbulence, owing to a more robust post-Great Financial Crisis (GFC) regulatory and supervisory framework, the credit quality of underlying portfolios and the alignment of banks' and investors' incentives. The main identified risks – procyclical dynamics from banks' reliance on NBFIs for credit protection, leverage and liquidity vulnerabilities among investors, and spillovers from bank-NBFI linkages – are still modest. However, this may change as SRT issuance activity grows, structures and risk transfer chains become more complex and banks rely more heavily on NBFIs for credit protection. In a stress scenario, SRTs may act as a transmission amplifier by tightening bank lending capacity and creating additional bank-NBFI adverse feedback loops.

Third, limited public disclosure and gaps in cross-sector and cross-border data – especially on investor funding structures, leverage and interlinkages – raise the risk that SRT-related vulnerabilities build up unnoticed. This underscores the need for more transparency and enhanced monitoring of these risks for individual banks and from a system-wide perspective. Regulators and supervisors should also guard against banks having to slow lending or let capital ratios fall if they cannot issue SRTs.

The rest of the article is structured as follows. The first section provides a brief introduction to SRTs and motivations for their use. The second examines global SRT trends and drivers. The third section analyses risks from a system-wide perspective. The article concludes with policy implications stemming from the findings.

The economics of SRTs

An SRT is a type of securitisation used by banks to transfer part of the credit risk associated with a pool of financial assets (typically loans) to investors. Those assets stay on the bank's balance sheet because the risk transfer takes place through a credit derivative or a financial guarantee. This arrangement differs from a cash or "true sale" securitisation, in which those assets are removed from the bank's balance sheet.

SRTs go by different names across jurisdictions. The Basel Framework refers to them as "synthetic securitisations", while other industry terms include "on-balance sheet securitisations" and "credit risk-sharing trades". In the EU and United Kingdom, the term "significant risk transfer" refers to both cash and synthetic securitisations that result in regulatory capital relief; in practice, most significant risk transfers are SRTs (ESRB (2025)). In the United States, the terms "capital relief trade" or "credit risk transfer" are also used.[2]

SRTs form part of banks' toolbox for managing their capital adequacy. At its most basic, increasing a bank's capital adequacy (CET1) ratio involves either increasing capital or reducing risk-weighted assets (RWA). Options for increasing capital include increasing retained earnings (through higher profitability or lower dividend distributions and share buybacks) or raising new equity. Options for reducing RWA include selling loans (directly or through true sale securitisations), hedging (using instruments such as SRTs, credit default swaps (CDS) or guarantees) or shifting to a lower-risk asset mix. In this context, SRTs are sometimes an easier and more cost-efficient solution than increasing capital, shrinking the balance sheet or selling (often illiquid) loans below their face value. They can also offer a more targeted approach than buying an index CDS on a public credit pool or obtaining credit insurance for an entire portfolio. By using SRTs, banks get the capital benefit as if they had sold the loan portfolio, but they retain ownership of the customer relationships.

SRTs are not used solely for regulatory capital relief. For example, banks may use SRTs for credit risk management by freeing up counterparty limits or reducing concentrations to particular industries, geographies or borrowers. Redeploying capital and enhancing lending capacity also enable banks to improve profitability.[3]

SRTs take three main forms (Graph 1). The first is a credit-linked note (CLN) on a reference pool of assets issued directly by the bank. The second is a financial guarantee or CDS on a portion of those assets between the bank and an investor. The third is a CLN issued by a special purpose vehicle (SPV) that has provided a guarantee or written CDS protection to the bank on the specific pool of assets. The issuance of a CLN, whether directly by the bank or indirectly by the SPV, allows the instrument to be sold more easily or pledged as collateral for repo financing (Cadwalader et al (2024); Parker (2025)).

Most SRTs today are funded or secured by financial collateral. The proceeds from the sale of the notes (for SPV-issued CLNs) or cash and high-quality securities collateral (for CDS and guarantees) are typically held by the bank itself or by a third-party custodian to cover any credit losses on an ongoing basis. This means that the bank has no counterparty risk to investor(s) or associated capital requirements.[4]

[2] Credit risk transfer is the name given in the United States to securities issued by government-sponsored enterprises (GSEs) to synthetically transfer a portion of mortgage credit risk to investors.

[3] It is difficult to assess how the capital freed up by SRTs is used by banks. Anecdotal evidence suggests that it has been reallocated to higher-return businesses, used to increase loans in the same business line or used to maintain a bank's capital ratio (eg after an acquisition). Klein et al (2025) find that banks use SRTs to shift their lending toward riskier green corporate loans. Osberghaus and Schepens (2025) find that euro area banks using SRTs have redirected freed up capital towards new lending.

[4] There is no collateral in unfunded transactions (other than potentially for margin purposes, eg for a bank CDS with an eligible guarantor). The risk weight of the tranched assets is replaced by the risk weight of the guarantor, such as a supranational institution or highly rated credit risk insurer.

Mechanics of capital relief for a typical three-tranche SRT structure[1] Graph 1

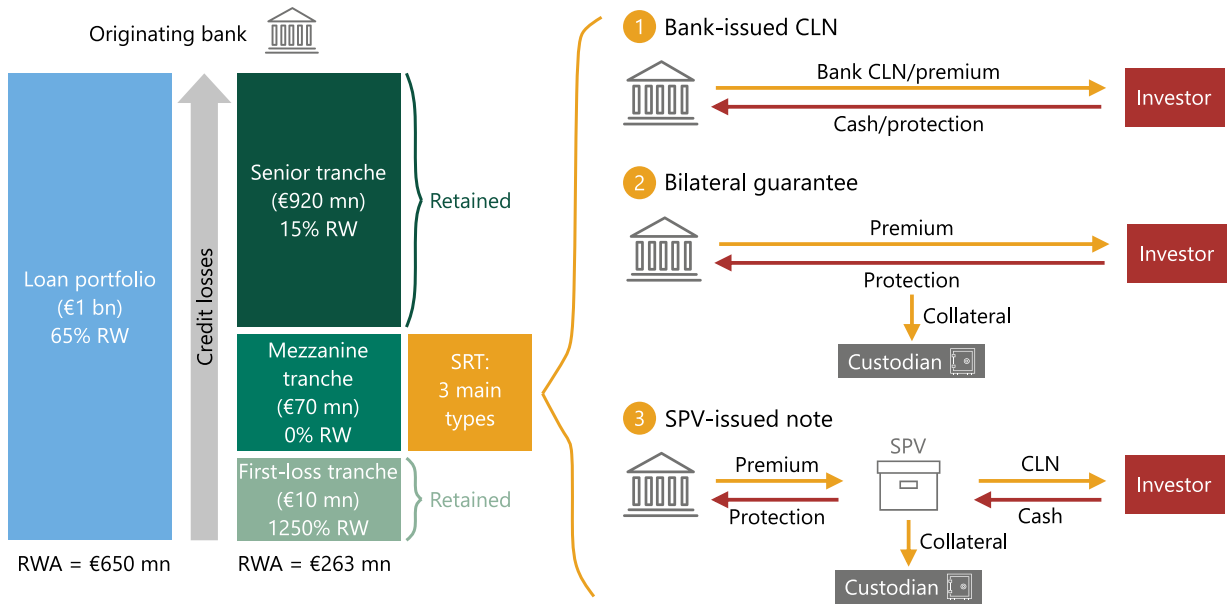

CLN = credit-linked note; RW = risk weight; RWA = risk-weighted assets; SPV = special purpose vehicle; SRT = synthetic risk transfer.

[1] See Annex A for the worked-out example.

Source: Authors' elaboration.

The originating bank pays a premium for credit protection. In a two-tranche SRT, it keeps the senior tranche and transfers the junior (first-loss) tranche to investors. In a three-tranche SRT, it transfers the mezzanine and sometimes the junior tranche. The return that the investor earns is the premium on credit protection plus the (floating rate) market return on the collateral. In the event of credit losses on the reference portfolio, the bank will write down the principal amount (for a CLN) or use the collateral to cover any losses, starting with the junior tranche. The bank is only exposed to loss when the protected tranche (in a two-tranche structure) is used up.

The economic appeal of SRTs for banks hinges on the trade-off between the cost of protection and the regulatory capital relief benefit. Banks have an incentive to economise on capital so that they can increase their profitability for shareholders. They therefore prefer to transfer credit risk for loans whose regulatory capital requirements are deemed to be higher than their own assessment of riskiness (Osberghaus and Schepens (2025)). In such cases, the reduction in RWA – and the associated savings in required capital – can outweigh the protection premium paid to investors. Graph 1 shows how a bank, by obtaining protection on the mezzanine tranche, is able to reduce its RWA for a loan portfolio by almost 60% (from €650 million to €263 million). Whether the transaction also enhances the bank's return on RWA will depend on the loan portfolio's yield vis-à-vis the cost of protection. Annex A includes an example of an SRT transaction providing capital relief and enhancing RWA returns.

SRTs appeal to specialised credit market investors because of their risk-adjusted returns and diversification benefits. Investors are attracted by the opportunity to access generally high-quality, bank-originated loans they may not otherwise obtain (eg exposure to smaller borrowers that are unrated and do not issue in public markets). Investors are also attracted by operational efficiencies from reliance on the

bank for loan origination, servicing and (if needed) workouts. Bank-issued CLNs are general unsecured obligations, so investors in these instruments are also exposed to the bank's counterparty risk, which they need to manage.[5]

The use of leverage can enhance investor returns. SRT investments have embedded leverage, ie a higher amount of exposure per unit of invested capital, because the loan portfolio is sliced into tranches that differ in payment priority. This allows for the amplification of risk and return for investors. Some SRT investors also use financial leverage to further enhance their returns. This typically takes one of two forms: a repo by pledging the CLN as collateral, or borrowing against other investor assets (fund financing). In the former case, the investor borrows a fraction of the CLN's value as determined by the repo haircut. In the latter case, the investor typically borrows against the fund's underlying portfolio (net asset value (NAV) loans) or against investor commitments for subscription line or capital call facilities. The incentives to use financial leverage depend in part on the investor's target return and the SRT structure. When investors have absolute return hurdles (as in the case of hedge funds and some private credit funds) and the structure comprises high-quality assets, then investors may need to use leverage to achieve their return hurdle.

Market overview and trends

SRTs have existed for a long time, but they gained prominence after the GFC. These instruments first appeared in the 1990s with the emergence of bespoke credit derivatives such as CDS. The recognition of these instruments as credit risk mitigants in Basel II gave further impetus to their growth in the 2000s (Renault (2025)). The GFC caused issuance to collapse and eroded confidence in synthetic structures and securitisations (FSB (2025a)). SRTs re-emerged post-GFC, underpinned by prudential reforms as well as macroeconomic and market developments that facilitated broader adoption. These included: Basel III, which incentivised SRT use given generally higher risk-based capital requirements for banks; changes to jurisdictional regulatory and supervisory frameworks that supported SRT market activity;[6] depressed bank equity valuations post-GFC, which enhanced the attractiveness of SRTs compared with issuing equity; and the growth of NBFIs and search for yield that led more investors – especially private credit funds and asset managers – to enter the market.

Monitoring developments in the global SRT market is challenging due to the scarcity of public information. Currently, there is no data repository or consistent regulatory reporting on issuance, pricing or credit performance across jurisdictions. Below, we combine information from various sources (bank reports, market surveys, commercial data providers and industry outreach) to provide a market overview.

The global SRT market has expanded fivefold since the mid-2010s. According to the International Association of Credit Portfolio Managers (IACPM (2025)), the annual issuance of SRT tranches grew from less than €5 billion in 2016 to €21 billion in 2024.

[5] Investors mitigate the counterparty risk in SRTs issued by smaller banks through a third-party custodian or segregation of cash collateral, or through provisions for the collateral to be moved to another bank if the issuing bank's credit rating deteriorates. Anecdotal evidence suggests that such practices have protected investors in the few cases of SRT-issuing bank failures to date.

[6] These include, for example, SRT guidelines by the European Banking Authority (EBA) and the extension of the "simple, transparent and standardised" (STS) framework to synthetic securitisations in the EU in 2021 and the issuance in 2023 of a Federal Reserve FAQ on using CLNs as a risk mitigant.

The tranches issued in 2024 provided protection to underlying loan portfolios amounting to €260 billion, compared with €55 billion in 2016 (Graph 2.A). Given an average SRT duration of three to five years, the total loan amount protected by all outstanding SRTs as of end-2024 was over €700 billion (IACPM (2025)). Estimates in this article, based on Pillar 3 disclosures, suggest that outstanding SRT loans amounted to almost €800 billion as of end-2024 (Graph 3.A), while market intelligence suggests continued growth in issuance in 2025.

Notwithstanding this growth, the global SRT market remains relatively small and concentrated in a few advanced economies. Assets protected by SRTs represented around 2% or less of total bank loans, or about 1% of consolidated total bank assets (BCBS (2026)), in the EU, United States, United Kingdom and Canada as of end-2024. The total loan amount protected is only a small proportion of the estimated global cash securitisation market (FSB (2025a)), though SRTs are more important in Europe (González and Triandafil (2023)).[7] SRT market characteristics also differ between the United States and other jurisdictions such as the EU, reflecting distinct financial system structures and regulatory frameworks (Annex C).

Corporate loans remain by far the largest asset class in SRTs, though the range of protected loan types has expanded recently. In 2024, corporate and small and medium-sized enterprise (SME) loans accounted for almost 80% of total issuance (purple bars in Graph 2.B). These loans are a core part of bank portfolios, but they typically generate low returns as most bank income from these borrowers comes from ancillary products and services. Despite their generally good credit quality, banks are incentivised to include these loans in SRTs due to their fairly high capital charges. Over time, the range of loan types in reference portfolios has expanded to include auto loans and residential mortgages (retail finance), as well as capital call facilities, commercial real estate (CRE) and leveraged finance. This trend was supported by the entry of smaller banks[8] and more US banks[9] into the market.

The range of banks issuing SRTs has become increasingly diverse over time. SRT issuance was historically dominated by a few, mostly global systemically important banks (G-SIBs) in the EU and United Kingdom. Around eight new banks have entered the market annually since 2016, raising the total number of issuers above 100 (Graph 2.C). These banks are European and, to a lesser extent, North American.

[7] According to the Association for Financial Markets in Europe (AFME) and IACPM, the underlying loan pool protected by SRTs that were issued in 2024 (€188 billion) exceeded the size of all true sale securitisations (€144 billion) in Europe – comprising the EU, United Kingdom and other European countries – if the retained parts of those securitisations (ie not placed in the market) are excluded.

[8] Smaller banks generally use more basic approaches to estimate their capital requirements. According to IACPM (2025), the share of loans accounted for by such banks has increased from less than 5% in 2016 to 20% in 2024, which shifts the asset profile of SRTs towards more retail exposures.

[9] The relatively higher – compared with Europe – risk weights for certain asset types in the United States (eg capital call facilities, high-quality auto and residential mortgage loans) make them more amenable to SRT inclusion.

Synthetic securitisation has grown and become more diverse over time

Graph 2

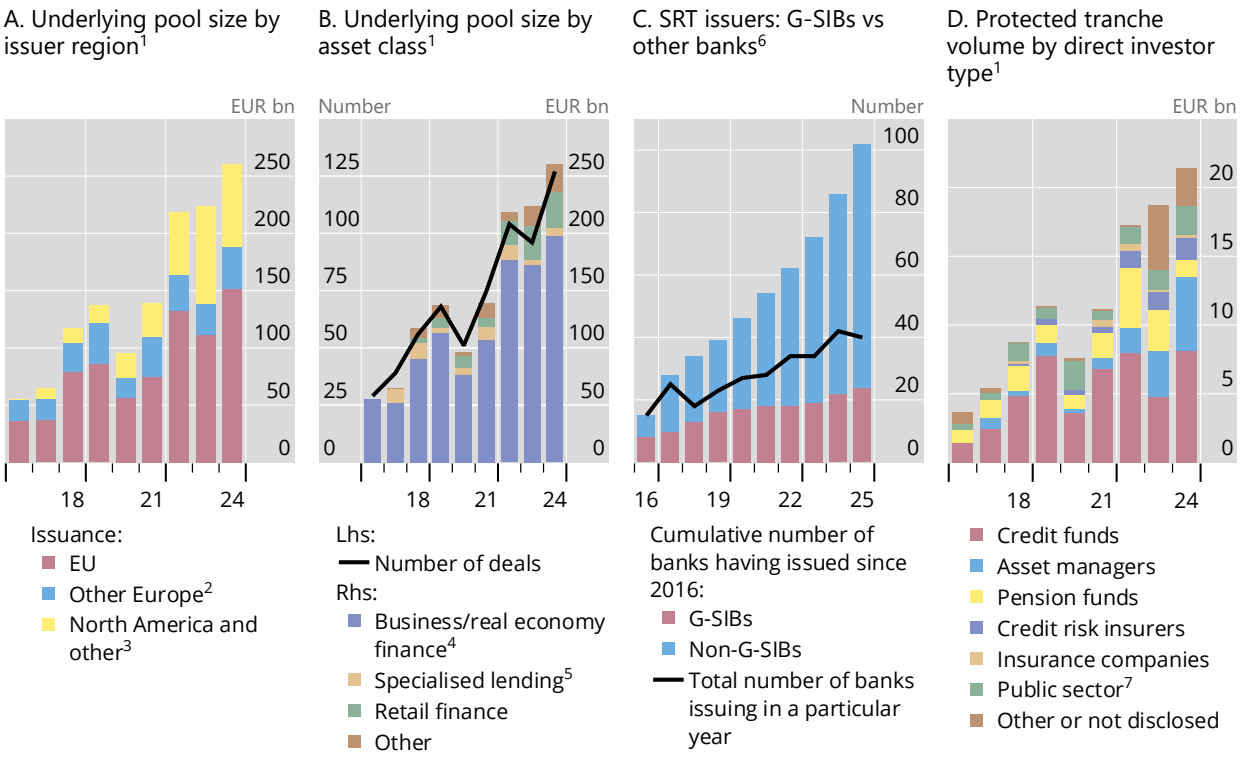

G-SIB = global systemically important bank; SRT = synthetic risk transfer.

[1] For a sample of 51 banks included in the IACPM's *Global SRT Bank Survey 2016–2024*. Volumes at inception. [2] Includes United Kingdom and Switzerland. [3] Includes South/Latin America, Asia, multi-country (same region) and unknown. [4] Includes corporate, small and medium-sized enterprise (SME) and trade finance loans. [5] Includes project/object finance and commercial real estate lending. [6] For the sample of banks available in the SCI database from 2016 onwards (excluding multilateral development banks that are SRT issuers). The term "bank" includes regulated credit institutions. Different issuing entities belonging to the same banking group are grouped together. The G-SIB category includes all banks that are currently or were previously designated as G-SIBs. [7] Includes central governments, central banks, international organisations and multilateral development banks.

Sources: IACPM (2025); Structured Credit Investor (SCI); authors' calculations.

Notwithstanding the broader range of issuing banks, the SRT market remains concentrated in terms of region, issuer and loan type. This is confirmed by analysing Pillar 3 disclosures of 44 SRT-issuing banks from Canada, Europe, Japan and the United Kingdom as of end-2024.[10] First, EU-domiciled banks represent around half of the global outstanding SRT amount (Graph 3.A). Canadian, UK and US banks account for a sizeable portion of the market, while Japanese banks are less frequent issuers. Second, the top 10 issuers represent 64% of the total outstanding amount. Third, 90% of the protected assets in the sample are wholesale (largely corporate) loans.

A bank's asset size and capital adequacy help explain which banks issue SRTs. Annex B includes an analysis of 28 EU banks that issued SRTs (based on Pillar 3 disclosures as of end-2024) vis-à-vis their peers. A simple regression analysis examining the likelihood of a bank issuing SRTs, based on key fundamentals, highlights two important factors: bank size (total assets) and the CET1 ratio. Specifically, larger banks are more likely to issue SRTs since they have greater resources and expertise and can leverage economies of scale in structured finance. In

[10] Pillar 3 disclosures for US banks do not provide sufficient information on SRTs to be able to estimate the extent of their use for each bank or to compare them with banks from other jurisdictions.

EU banks dominate the SRT market, while SRT use varies widely across banks[1]

As of end-2024[2] Graph 3

A. Estimated outstanding SRT-protected loans for banks headquartered in each jurisdiction[3]

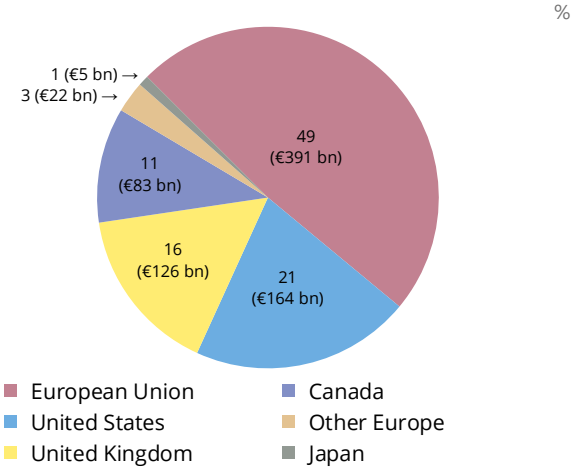

B. Estimated capital relief and share of loan portfolio protected for SRT-issuing G-SIBs and other banks[4]

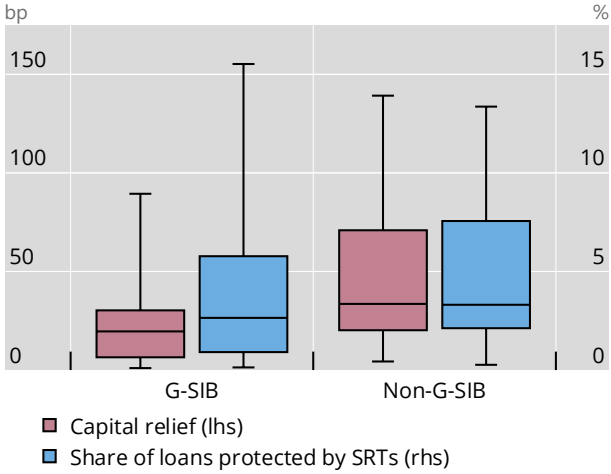

G-SIB = global systemically important bank; SRT = synthetic risk transfer.

[1] Based on Pillar 3 disclosures of five Canadian, 33 European, one Japanese and five UK banks that issued SRTs. The proportion of the SRT loan portfolio that is protected (junior and/or mezzanine tranche sold to investors) is assumed to be 10%. [2] As of end-January 2025 for Canadian banks and end-March 2025 for Japanese banks due to country-specific Pillar 3 reporting conventions. [3] The shares do not add up to 100 due to rounding. The figure for the United States is based on Brizzi et al (2025). [4] The box-and-whisker plots show the minimum, first-quartile, median, third-quartile and maximum values of the sample of all SRT-issuing banks. To arrive at the estimated basis points of CET1 capital relief, the average riskiness (RWA to assets) for each bank's SRT loan portfolio (pre-deal) is assumed to be the higher of 50% or that bank's average riskiness for the specific loan type and capital calculation approach.

Sources: Brizzi et al (2025); banks' Pillar 3 disclosures; S&P Global Market Intelligence; authors' calculations.

addition, banks with lower capital adequacy levels are more likely to issue SRTs. For an EU bank with an 19% CET1 ratio (average of the sample), a 1 percentage point decrease in that ratio increases the likelihood that it will issue SRTs by 25 basis points.

Only a few banks appear to be materially reliant on SRTs in terms of capital relief or the proportion of the loan portfolio protected (Graph 3.B). The average amount of capital relief for the sample of 44 SRT-issuing banks is estimated to be around 43 basis points of CET1 capital, though it exceeds 100 basis points in a few cases. This is a small portion of banks' CET1 ratio compared with a sector-wide average level of 14–16% in Canada, the EU and United Kingdom. Similarly, the average loan portfolio protected by SRTs for the 44 SRT-issuing banks is estimated to be only around 5% of their total loans, though it exceeds 10% for a few banks. The outliers are non-G-SIBs, even though in absolute terms G-SIBs account for half of the outstanding SRT amount.[11] These outlier banks are not small: most are designated as domestic systemically important in their home jurisdiction. These banks may rely more on SRTs because of more limited access to capital markets and other credit portfolio management tools given their size and credit rating. Their smaller size also means that a typical SRT transaction (around €2 billion) tends to have greater capital impact.

Although the SRT investor base has expanded and become more diverse over time, it remains concentrated (Graph 2.D). Prior to the Covid-19 pandemic, a few

[11] Anecdotal evidence suggests a similar trend in the United States, with some regional banks being most active (in terms of capital relief or relative to the size of their balance sheet) in the SRT market.

specialised credit funds and pension funds accounted for three fourths of annual SRT investments. However, their share has declined more recently as private credit funds, hedge funds and large multi-strategy asset managers have entered the market.[12] Supranational institutions currently represent around 10% of the investor base, with the aim of promoting SRT issuance by banks in less advanced economies or supporting specific activities (eg green or SME lending). Insurers account for less than 10% of SRT investments, offering mostly unfunded protection (liability side of their balance sheets).[13] However, the total number of SRT investors is still limited and the market is concentrated (Cortes et al (2025)). In the United Kingdom, for example, the top 10 protection sellers accounted for nearly 60% of SRT volumes in recent years.

Some factors suggest that the SRT market will continue growing. More banks are turning to this instrument to optimise their balance sheet, enhance profitability and manage credit risk. New economic priorities (eg digitalisation, climate transition and defence) may raise loan demand and incentivise banks to make greater use of SRTs in some jurisdictions. Regulatory initiatives in the EU appear supportive of synthetic securitisation (Coelho et al (2025)). The growth of private credit funds and multi-strategy asset managers provides additional pools of capital to invest in this market. Back-of-the-envelope calculations suggest scope for further growth (Box A).

Box A

Assessing the potential for the SRT market to grow in the future

One way to assess the potential future size of the SRT market is to assume that the entire banking sector issues SRTs in the same proportion as the "average" SRT-issuing bank. Taking the EU as an example, the average SRT-issuing bank had around 5% of its loan portfolio protected by SRTs (based on Pillar 3 disclosures), corresponding to almost 50 basis points of CET1 capital relief. For EU banks, the outstanding loan portfolio protected by SRTs was almost €400 billion as of end-2024, corresponding to 2.1% of the total loan portfolio for the banking sector. Assuming that SRT use across the sector reaches the same level as that of the average SRT-issuing bank results in an outstanding amount of more than €800 billion, about double the current figure. The order of magnitude of potential growth would be even greater in other, less mature markets.

In practice, the use of SRTs is unlikely to reach this figure given differences in bank size, loan portfolio and business strategy, investor preferences and regulatory stance. In particular, two regulatory and supervisory constraints may become relevant for some banks before reaching that figure. First, the leverage ratio acts as a non-risk-based structural backstop for SRT growth. While SRTs reduce RWA and improve risk-based capital ratios, they do not lower leverage exposure as the underlying loans remain on the balance sheet and continue to be subject to the leverage ratio. Second, some supervisory authorities may not be comfortable with banks relying on SRTs to such an extent that, if the SRT market were to shut down, some banks would drop below their minimum capital adequacy requirements unless they cut back on lending or raise equity. The capital relief allowed by supervisors would in such case depend on a bank's capital headroom above regulatory requirements.

[12] Market intelligence suggests that pension funds, along with sovereign wealth funds and endowments, are the main indirect investors in SRTs through their participation in specialised credit funds and diversified asset managers.

[13] The relatively short duration given the types of loan protected and the lack of an external credit rating of most SRTs reduces their appeal as an investment product for insurers.

SRT risks: a system-wide perspective

SRTs are generally better structured and subject to a more robust regulatory and supervisory framework than they were before the GFC. The reforms adopted following the 2008 crisis included better calibration of capital requirements for securitisation exposures, enhanced supervisory reviews, minimum risk retention requirements, reduced reliance on external credit ratings, stronger disclosures, and the discouragement or ban of re-securitisations (Coelho et al (2025); FSB (2025a); BCBS (2026)). These reforms have helped to align the incentives of banks and investors and ensure true risk transfer, materially strengthening the microprudential underpinnings of SRTs.

Against this background, the remainder of this section sets out emerging risks in the SRT market from a system-wide perspective. While these risks appear limited at present given the market's size and participant profiles, they warrant attention as issuance grows, structures evolve and bank-NBFI interlinkages deepen.

Rollover risk and procyclicality

The most direct source of risk stemming from SRTs is rollover (or flowback) risk for the issuing bank. Market intelligence suggests that banks generally structure SRT maturities to match or exceed the maturity profile of the underlying loans. Yet even when maturities are aligned, rollover risk can arise if the bank intends to maintain or expand lending beyond the SRT's maturity date and therefore needs to replace expiring protection (ESRB (2025)). Rollover risk tends to be more acute for banks with large, one-off issuances or where multiple transactions mature in close succession. The concentrated nature of the SRT investor base amplifies this risk, as changes in risk appetite or funding conditions among a small group of investors can have outsized effects on issuance volumes and pricing.

Rollover risk implies that banks become vulnerable to shifts in investors' ability and willingness to provide new protection, potentially giving rise to greater procyclicality of credit supply. Banks may find such protection expensive or unavailable precisely when capital needs are greatest.[14] A failure to extend protection would force banks to reduce lending growth or let their capital adequacy ratio decline. Tying banks' lending capacity to the SRT investor base suggests that, in the event of a broad economic downturn, the aggregate credit supply could become more procyclical and exacerbate the downturn (BCBS (2026)).

Such a risk has not materialised to date even during episodes of market turbulence. This can be attributed to the strong credit performance in underlying loan portfolios (EBA (2020)),[15] the limited reliance so far on SRTs for managing banks' CET1

[14] Credit rating agencies also monitor banks' reliance on SRTs. For example, Moody's (2025) notes that a capital improvement of more than 100 basis points may reflect overreliance on SRTs and other risk transfer transactions, since they provide materially less protection to bank creditors than equity.

[15] There is no evidence to date of weakening underwriting standards due to risk transfer reducing banks' incentives to monitor loans properly. Osberghaus and Schepens (2025) argue that banks using SRTs reduce monitoring relative to other banks that are lending to the same firm. However, the fact that the bank continues to hold the senior tranche and manage the loan, and that it retains an ongoing relationship with the protection provider(s), should help mitigate any misalignment of incentives. Moreover, some jurisdictions (eg EU) require banks to keep an economic interest in the securitised portfolio ("skin in the game") and to not differentiate lending standards by portfolio.

ratios and the absence of a "maturity wall" in which many SRTs mature at around the same time (EBA (2025)). The presence of dedicated funds that act as stabilising buyers in times of stress (such as during Covid-19) may also help mitigate this risk. Nevertheless, some supervisory authorities have highlighted the importance of incorporating flowback risk into bank capital planning frameworks (PRA (2025)) or have developed methodologies to assess and mitigate it (Finansinspektionen (2021)).

Investor leverage and liquidity risks

Leverage can lead to forced asset sales in times of stress, but the use of leverage in SRT investor portfolios appears to be modest.[16] While public information on SRT investor leverage is unavailable, anecdotal evidence suggests that it remains limited. Banks confine SRT repo financing to a subset of investors and typically subject such financing to conservative safeguards: haircuts of 40–60%, daily margining and cross-collateralisation requirements that allow lenders to draw on other assets if needed.[17] The seniority and over-collateralisation of financed positions, together with the lower mark-to-market volatility of SRT tranches compared with publicly traded structured credit, likely provide a meaningful buffer before losses are realised. Moreover, the bank providing the financing to the investor is subject to capital requirements for that exposure. Market intelligence suggests that SRT-financing banks are generally distinct from those issuing the instruments, though the fact that many of these banks are located in other jurisdictions complicates the monitoring of these exposures.

Anecdotal evidence also indicates that liquidity mismatch among SRT investors is limited. SRT investments are generally held until maturity owing to their bespoke nature, small market size and, in some cases, transfer restrictions and confidentiality provisions tied to the underlying portfolios. Hence, in principle, liquidity mismatch in open-ended funds investing in SRTs may create run dynamics in times of stress (ESRB (2025)). However, this risk appears contained in practice: many SRT funds use closed-end structures that allow them to manage redemptions without having to sell holdings. Most open-ended funds use an evergreen structure less exposed to disruptive redemptions,[18] allocate a small fraction of their multi-asset portfolios to SRTs or manage redemptions through bank lines and internal liquidity buffers.

Interlinkages and risk transfer chains

The growing and increasingly complex risk transfer chains among SRT market participants can create hidden vulnerabilities (Table 1, Graph 4). These chains stem

[16] Differences in SRT market characteristics across jurisdictions may also affect the use of leverage (Annex C). For example, US transactions typically have thicker tranches (ie a greater proportion of the loan portfolio needs to be sold to investors) in order to maximise capital relief given the capital rules in the United States, which incentivises greater use of leverage by investors to enhance returns.

[17] Drawing on Seer Capital (2025), if one assumes that only part of the SRT investor base (ie asset managers and credit/hedge funds) are potential users of leverage, that half of them actually use leverage and that users generally borrow around half of the SRT amount they purchase (given the haircuts), then the total leverage outstanding in the SRT market would be around €10 billion.

[18] An evergreen fund has no fixed end date and is designed to recycle capital on a continuous basis. Investors have the ability to enter or exit on preset schedules rather than wait for fund life cycles.

Illustrative typology of risk transfer chains involving SRTs[1]

Table 1

Risk transfer chain	What it involves	Why it matters for financial stability
1) Bank and insurers *(bank → insurer)*	Unfunded SRTs where insurers provide credit protection.	Exposes banks to risk in insurance sector and enforceability of guarantees under stress.
2) Banking system *(bank → funds → (different) bank)*	Funds sell protection but finance positions via repo or credit lines from other banks.	Creates circles of risk and reintroduces exposure to the banking system if losses exceed protected tranches; amplifies leverage and liquidity risk.
3) Non-bank sectors *(fund → bank → same fund)*	A fund sells SRT protection to a bank on a portfolio that includes exposures to that fund complex or to companies the fund owns or finances.	Creates correlation and conflict-of-interest risk; losses on underlying assets coincide with SRT obligations.
4) Multi/SRT repack *(bank → fund → bank → insurer)*	Complex structures can arise when a bank transfers risk to a fund, the fund finances its position through another bank and an insurer provides credit protection to the financing vehicle.	Adds structural complexity and opacity; concentrates counterparty risk across banks, funds and insurers; amplifies contagion if one link fails.
5) Re-securitisation of SRTs	Repackaging of SRT exposures into new tranched structures, spreading risk across the financial system.	Potential to recreate pre-GFC-style opacity and amplification; systemic contagion risk.
6) Cross-border spillovers	Issuers (and their loan portfolios protected by SRTs), investors and financing providers operate across multiple jurisdictions.	Transmission of shocks globally; regulatory arbitrage and coordination challenges.

SRT = synthetic risk transfer.

[1] See Graph 4 for a visual illustration of the various risk transfer chains.

Source: Authors' elaboration.

SRTs create additional linkages between banks and NBFIs[1]

Graph 4

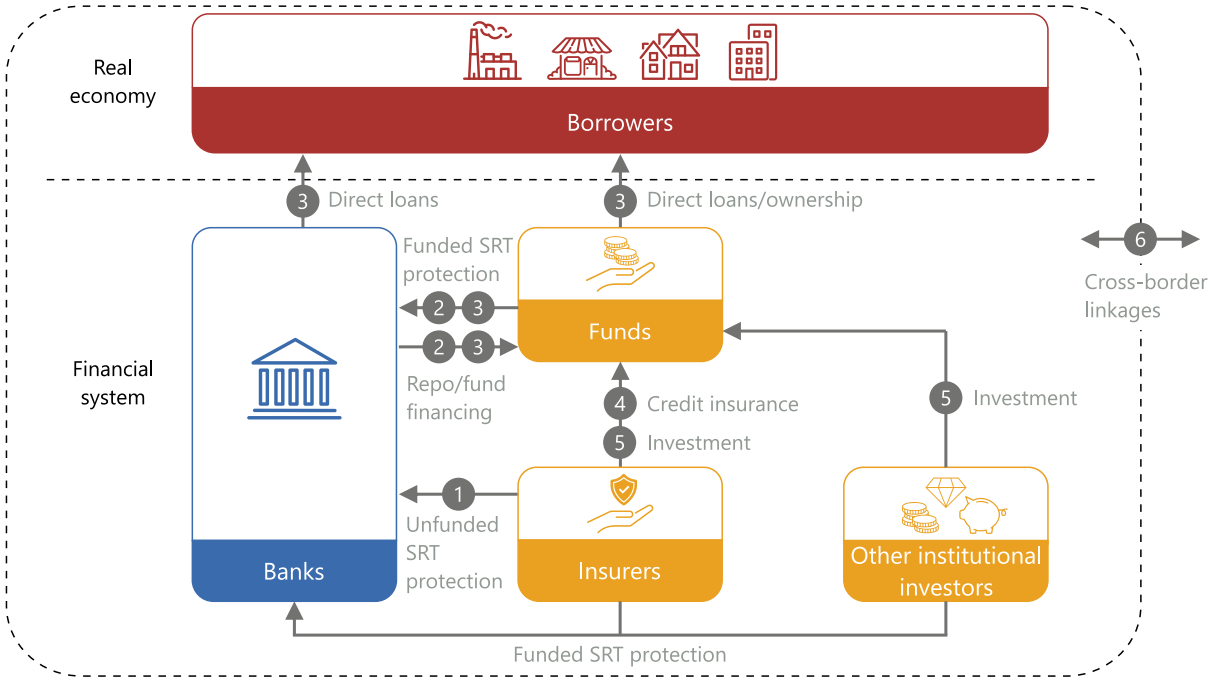

NBFIs =non-bank financial institutions; SRT = synthetic risk transfer.

[1] See Table 1 for a description of the various risk transfer chains. The arrows only show linkages between entities relating to SRTs. Fund financing includes net asset value (NAV) loans and capital call facilities. Other institutional investors include pension funds, sovereign wealth funds and endowments.

Source: Authors' elaboration.

from the expanding set of bank issuers and investors, the broader range of financing structures and the diversity of loan types included in SRT portfolios. Linkages across banks, funds and other institutional investors take many forms and continue to evolve, reflecting the wider growth of bank-NBFI interconnectedness documented in recent work (Acharya et al (2024)).

SRTs can create additional transmission channels during stress. One example is "circles of risk" (ESRB (2025)), where credit risk transferred from banks to funds via SRTs can return indirectly to the banking sector as a result of other banks financing the purchases of those investors (type 2 shown in Table 1 and Graph 4). In some transactions, the reference portfolio may include loans to firms that the fund owns or finances or the provision of capital call facilities to the private equity affiliates of those funds (type 3). This raises the risk of credit exposures ostensibly transferred out of the banking sector being indirectly reabsorbed. Repackaging of SRT exposures to other investors (types 4 and 5) adds further layers of complexity and can, in principle, recreate pre-GFC-type amplification mechanisms, although evidence to date of such structures and transmission channels is limited.

The transfer of credit risk across borders through SRTs (type 6) can enhance risk diversification, but it can also increase the transmission of shocks globally. For example, European and Japanese banks reportedly transfer credit risk through SRTs to US-based hedge funds and private credit funds while concurrently providing financing to some of those entities, creating the potential for negative feedback loops. Anecdotal evidence suggests that the scale of these loops remains modest.

Limited and fragmented information on interlinkages and risk transfer chains heightens the potential for SRT-related risks to build undetected. Authorities in many jurisdictions require notification or reserve the right to review SRTs as part of the supervisory process. However, opacity around investor characteristics, funding links and cross-border exposures limits authorities' ability to monitor concentrations, leverage and interconnectedness. This increases the likelihood that vulnerabilities remain unnoticed until stress materialises. In a hypothetical stress scenario, while SRTs would not be the primary shock driver, they could contribute to a procyclical tightening of bank lending capacity and enhance bank-NBFI adverse feedback loops.

Conclusion

SRTs are becoming a significant element of the evolving credit risk-sharing ecosystem between banks and NBFIs. The growth of SRTs carries both benefits and risks for financial stability. If done in a prudent manner, risk transfer can contribute to a more diverse and robust financing ecosystem. However, this depends on whether investors are well placed to assume the risk given their investment horizon, funding structure, sensitivity to mark-to-market dynamics and ability to withstand losses in stress.

While the SRT market has been resilient in recent episodes of market turbulence, it has not yet experienced a prolonged credit market downturn. The main risk channels identified in this article – procyclical dynamics from banks' reliance on NBFIs for credit protection, leverage and liquidity vulnerabilities among investors and spillovers from bank-NBFI linkages – are still modest and evolving, but they highlight how SRTs can contribute to the transmission of shocks across the system.

Three main policy implications emerge from the analysis, reflecting both system-wide and bank-level considerations. First, greater transparency is needed to

monitor SRT-related risks as the market continues to expand. This includes public disclosures of the capital impact of SRTs and enhanced information-sharing among authorities – including on a cross-border basis – on investor funding structures, leverage and interlinkages. Second, from a macroprudential perspective, the incorporation of SRTs into system-wide stress testing exercises could help shed light on bank-NBFI spillovers under severe but plausible scenarios. And third, at the microprudential level, regulatory and supervisory attention is needed to ensure that the expansion of SRTs does not expose banks to rollover risk and procyclical behaviour. The adoption of internationally consistent securitisation frameworks, together with policies to address NBFI vulnerabilities, would reinforce these safeguards.

Finally, this article raises some broader questions about the role of SRTs in the financial system. These include the extent to which SRTs support additional lending to the real economy; the degree to which their use is shaped by the interaction of capital rules, loan types and economic conditions; and their implications for the overall level and composition of capital in the banking sector. These issues fall beyond the scope of this analysis but merit further investigation.

References

Acharya, V, N Cetorelli and B Tuckman (2024): "Where do banks end and NBFIs begin?", *NBER Working Paper*, no 32316.

Basel Committee on Banking Supervision (BCBS) (2026): *Synthetic risk transfers*.

Brizzi, J, J Gerlach and J Kelly (2025): "Banking trends: synthetic risk transfers", Federal Reserve Bank of Philadelphia, *Economic Insights*, vol 10, no 3.

Cadwalader, KPMG and TWK Advisors LLC (2024): *Credit risk transfers (CRTs) – A handbook for US banks*.

Coelho, R, R Corrias and J-P Svoronos (2025): "Post-GFC securitisation reforms and new initiatives: a comparative analysis", *FSI Insights on policy implementation*, no 71.

Cortes, F, G Fernandez Dionis, Y Li, S Ramirez and X Zhang (2025): "Recycling risk: synthetic risk transfers", *IMF Working Papers*, no WP/25/200.

European Banking Authority (EBA) (2020): *Report on STS framework for synthetic securitisation under Article 45 of Regulation (EU) 2017/2402*.

——— (2025): *Risk assessment report of the European Banking Authority*.

European Systemic Risk Board (ESRB) (2025): *Unveiling the impact of STS on-balance-sheet securitisation on EU financial stability*.

Financial Stability Board (FSB) (2025a): *Evaluation of the effects of the G20 financial regulatory reforms on securitisation – Final report*.

——— (2025b): *Enhancing the resilience of Nonbank financial intermediation – Progress report*.

Finansinspektionen (2021): *Updated Pillar 2 method for assessing flowback risk associated with securitisation*.

González, F and C Triandafil (2023): "The European significant risk transfer securitisation market", *ESRB Occasional Paper Series*, no 23.

International Association of Credit Portfolio Managers (IACPM) (2025): *Global SRT Bank Survey 2016–2024*.

Klein, P, A Nitschke and A Pfingsten (2025): *Financing the green transition: A case for synthetic capital relief trades*.

Moody's (2025): *Rating methodology: Banks*.

Osberghaus, A and G Schepens (2025): "Synthetic, but how much risk transfer?", unpublished working paper.

Parker, E (2025): "Synthetic risk transfer SRT in 2025", *Insights*, Mayer Brown.

Pennacchi, G and J Santos (2021): "Why do banks target ROE?", *Journal of Financial Stability*, vol 54, June.

Prudential Regulation Authority (PRA) (2025): *Significant risk transfer financing: Prudential expectations*.

Renault, O (2025): "Significant risk transfer (SRT) Chronicles 2025", Pemberton.

Seer Capital (2025): *Reg Cap leverage: clearing misconceptions*.

Annex A

This annex demonstrates how banks can use SRTs to reduce the risk weight of a loan portfolio, thereby improving their capital adequacy ratio and return (all other things equal). The calculations of RWA and net income have been simplified by excluding accounting, tax, amortisation, expected loss and capital redeployment considerations.

Impact on RWA and required capital. Consider Bank A, which holds a €1 billion portfolio of loans on its balance sheet, with a risk weight of 65%. To reduce its RWA, Bank A slices the portfolio into three tranches: the junior (first-loss) tranche, the mezzanine tranche and the senior tranche, with tranche sizes of 1%, 7% and 92%, respectively. Bank A retains the senior tranche, representing €920 million of the exposure, which carries a risk weight of 15% as per EU capital rules. The bank also retains the first-loss tranche (1% of portfolio), which is subject to a 1,250% risk weight assigned to highly risky junior securitisation tranches. The mezzanine tranche (7%) is hedged through a CLN sold to investors, reducing its risk weight to 0%. As a result, Bank A reduces its RWA from €650 billion to €263 million (Table A.1). Assuming a 12.5% CET1 ratio, the use of the SRT allows the bank to lower its required capital to €33 million (€263 million × 12.5%), compared with €82 million without the SRT.

Impact on income. Assume Bank A charges borrowers a 3% interest rate on its loan portfolio (net of funding costs), generating an annual income of €30 million. While net income is the same after implementing the SRT, Bank A hedges a portion of the credit risk at a premium of 7%, which costs €5 million. Consequently, its post-SRT net income decreases to €25 million. However, owing to the reduction in RWA, the bank's net income as a percentage of RWA rises from 5% pre-SRT to 10% post-SRT.

SRT illustrative example (estimates at inception) Table A.1

	Pre-SRT	**Post-SRT**			
	Ref portfolio	Senior tranche	Mezzanine tranche	First-loss tranche	Total
Size (€ mn)	1,000	920	70	10	1,000
Risk weight (%)	65	15	0	1,250	
RWA (€ mn)	650	138	0	125	263
CET1 (%)	12.5	12.5	0	12.5	
Required capital (€ mn)	82	17		16	33
Net interest rate on loans (%)	3				3
Net interest income (€ mn)	30				30
Protection premium (%)			7		
Cost of protection (%)			4.9		
Net income (€ mn)	30				25
Net income/RWA (%)	5				10

SRT = synthetic risk transfer.

Source: Authors' calculations.

Annex B

This Annex aims to shed light on the characteristics of banks that have issued an SRT. The analysis focuses on EU banks given the large sample and uniform disclosures and regulatory framework.[19] It highlights differences in financial indicators between banks that issued SRTs and those that did not in 2024. Several observations are worth noting.

First, banks that issued SRTs were generally larger in size. The median total assets of the 28 SRT-issuing banks exceeded €290 billion, which is more than four times the size of the 78 non-SRT-issuing banks. This is not surprising, as larger banks typically have greater resources and benefit from scale in structured finance. In addition, the fixed costs associated with issuing SRTs can be prohibitively high for smaller banks.

Second, on average, SRT-issuing banks tended to have lower CET1 ratios compared with their non-SRT-issuing counterparts (Graph B.1.A). All else being equal, banks with lower CET1 ratios would be expected to use SRTs to enhance their capital ratios.

Third, publicly listed SRT-issuing banks, on average, have a lower price-to-book ratio (PBR) compared with non-SRT-issuing banks (Graph B.1.B).[20] A lower PBR means that issuing common equity may dilute value for existing shareholders and could trigger an unfavourable market reaction. Consequently, banks with lower PBRs are more likely to reduce RWA using SRTs to help raise their CET1 ratio.

Fourth, SRT-issuing banks tend to have a higher return-on-equity (RoE) (Graph B.1.C). However, there is a causality issue to consider. On one hand, SRTs can enable banks to free up capital, allowing them to engage in more profitable activities and thereby enhancing their RoE. This is particularly relevant for banks that prioritise RoE over earnings in their performance targets (Pennacchi and Santos (2021)). On the other hand, investors are more inclined to provide credit protection to banks that are perceived to be more resilient because they report higher profitability.

Fifth, SRT-issuing banks exhibit higher loan concentration, defined as the ratio of corporate to total loans (Graph B.1.D). This aligns with market observations suggesting that banks use SRTs mainly for corporate loans to manage concentration risks and optimise their overall capital structure.

To assess how various factors influence the probability of a bank issuing SRTs, we estimate the following probit model:[21]

$$\text{Prob (SRT issuance}_i > 0) = \Phi(\text{constant} + \beta_i \times \text{financial strength variable}_i),$$

where $\Phi(\bullet)$ represents the standard normal cumulative distribution.

[19] The sample is drawn from the EBA's annual EU-wide transparency exercise, which covers 119 major banks across the EU. However, due to data availability and limitations, the sample size is reduced to 106 banks. This analysis relies on bank-level data from Pillar 3 disclosures, linking them with bank fundamentals. Notably, it does not explore how banks select loan portfolios for SRTs.

[20] Focusing on the PBR reduces the sample to 53 publicly listed banks, of which 24 issued SRTs.

[21] A probit model is a type of regression model used when the dependent variable (eg SRT issuance) is binary, ie it takes on two possible outcomes.

Bank fundamentals and SRT issuance[1]

Graph B.1

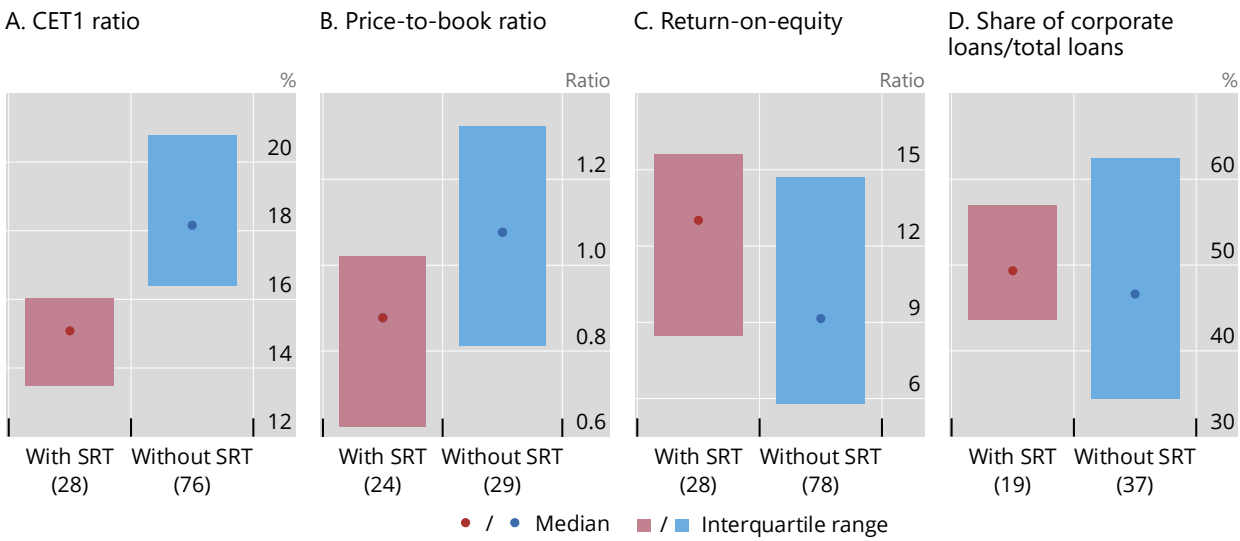

CET1 = common equity tier 1; SRT = synthetic risk transfer.

[1] End-2024 values for 106 European banks included in the EBA's EU-wide transparency exercise (2025). The sample size differs due to data limitations and the fact that some of these banks are not publicly listed; number of banks for each variable is shown in brackets.

Sources: European Banking Authority (EBA); banks' Pillar 3 disclosures; S&P Global Market Intelligence; authors' calculations.

Table B.1 presents key estimation results. The analysis begins with a general model (Probit 1) that includes all of the variables mentioned above, but RoE and loan concentration are found to be statistically insignificant. Further examination leads to a more parsimonious model (Probit 2) with only total assets and CET1 ratio being statistically significant in affecting the probability of a bank issuing SRTs. Specifically, for a bank with an average CET1 ratio, a 1-percentage point decrease in the CET1 ratio increases the probability that it will use SRTs by 0.25%, all else being equal.[22] Similarly, a larger bank in terms of asset size is more likely to issue SRTs. It is important to note that these estimations are based on cross-sectional data, which do not account for the dynamic nature of SRT issuance over time, and further research is needed to incorporate temporal dynamics into the analysis.

Probability of bank issuing SRTs[1]

Table B.1

	Constant	Total assets (€ trn)	HCET1[2] (%)	PBR (%)	Share of corporate loans (%)	RoE (%)	Pseudo R-squared[3]	No of obs
Probit 1	2.38 (2.06)	3.23* (1.81)	–0.21 (0.14)	–1.41 (1.52)	0.03 (0.02)	–0.04 (0.11)	0.47	37
Probit 2	1.97** (0.90)	1.86*** (0.63)	–0.25*** (0.08)				0.36	104

PBR = price-to-book ratio; SRT = synthetic risk transfer. Figures in brackets are standard errors; ***/**/* denotes statistical significance at the 1/5/10% level, respectively.

[1] End-2024 values for 106 European banks included in the EBA's EU-wide transparency exercise (2025). The sample size differs due to data limitations and the fact that some of these banks are not publicly listed. [2] Common equity tier 1 ratio minus the regulatory minimum of 4.5%. [3] Pseudo-R-squared is a goodness-of-fit measure based on maximum likelihood estimation.

Sources: European Banking Authority (EBA); banks' Pillar 3 disclosures; S&P Global Market Intelligence; authors' calculations.

[22] One standard deviation of the CET1 ratio is 8%.

Annex C

This Annex compares SRT market characteristics in the United States and other jurisdictions such as the EU (Table C.1). In the EU, given the bank-centric financial system, broad adoption of Basel II and III, and supportive regulatory and supervisory framework, banks have long used SRTs to transfer credit risk. The SRT markets in Canada and the United Kingdom are broadly similar to that of the EU in terms of regulatory capital treatment, structural features, loan portfolios and the investor base. In the United States, the existence of deep capital markets for corporate borrowers and the availability of government-supported SME loans may have reduced the need for SRTs for these asset classes, while differences in capital rules have resulted in distinct loan types protected by SRTs. Supervisory notification expectations, reporting requirements and public disclosures also differ across jurisdictions (BCBS (2026)).

The SRT market differs in the EU and the United States — Table C.1

SRT market characteristics	EU	US[1]
Estimated outstanding amount protected (end-2024)	€391 billion (2.1% of banking sector loans)	$170 billion (1.4% of banking sector loans)
Risk transfer instrument	Mostly bank-issued CLNs (SPV use has declined over time) Some use of unfunded credit insurance	Two thirds are bilateral CDS between the bank and a counterparty More bank-issued CLNs since issuance of FAQ by Federal Reserve Board in 2023
Typical structure	Three tranches, with a retained first-loss piece of 0.5–1.5% and a mezzanine tranche of 5.5–7.5% of the reference portfolio, respectively Two-tranche structure is also used occasionally	Two tranches, with a first-loss piece of 12.5% of the reference portfolio
Underlying loan portfolio composition	Corporates (75%), SMEs, commercial real estate, consumer and mortgage loans 19% of underlying loans are to non-EU borrowers	Corporates, capital call facilities, high-quality auto loans and residential mortgages
Investor profile	Credit funds, asset managers, pension funds, supranationals (eg European Investment Fund), non-life insurance companies (unfunded credit protection) Approximately one third of euro area SRTs involve foreign investors, predominantly US-based	Credit funds, asset managers, life insurance companies, pension funds
Regulatory and supervisory framework	Supervisory notification and approval (on a non-objection basis) for SRTs to obtain capital relief Mandatory risk retention and supervisory reporting requirements SRTs are eligible for STS treatment since 2021 (almost half of all outstanding EU SRTs are STS)	No mandatory notification or reporting requirement if protection is through CDS Bank-issued CLNs require supervisory approval for capital relief and are subject to bank-specific limits (lower of 100% of the bank holding company's total capital or $20 billion in reference loans)
Pillar 3 disclosures	Some specific reporting requirements on SRTs (eg exposures, RWA, loan type)	No specific reporting requirements on SRTs

CDS = credit default swap; CLN = credit-linked note; FAQ = frequently asked questions; RWA = risk-weighted assets; SME = small and medium-sized enterprises; SPV = special purpose vehicle; SRT = synthetic risk transfer; STS = simple, transparent, and standardised.

[1] US SRTs do not include mortgage credit risk transfer instruments issued by government-sponsored enterprises (GSEs) to investors.

Sources: BCBS (2026); Brizzi et al (2025); IACPM (2025); Board of Governors of the Federal Reserve System; Eurostat; banks' Pillar 3 disclosures; industry outreach; authors' elaboration.

Batuhan Koyuncu
koyuncu@cs.uni-saarland.de

Byeungchun Kwon
byeungchun.kwon@bis.org

Marco Lombardi
marco.lombardi@bis.org

Fernando Perez-Cruz
fernando.perez-cruz@bis.org

Hyun Song Shin
hyunsong.shin@bis.org

BISTRO: a general purpose oracle for macroeconomic time series[1]

Predictions of macroeconomic variables are a key input to economic policy, yet traditional econometric approaches have the limitation that the model needs to be tailored to the specific task. The advent of large language models (LLMs) opens up the tantalising prospect that a single general model can tackle a wide variety of tasks. This article introduces the BIS Time-series Regression Oracle (BISTRO), a general purpose time series model for macroeconomic forecasting. Building on the transformer architecture underlying LLMs, BISTRO is fine-tuned on the large repository of macroeconomic data maintained at the BIS. We put the model through its paces by assessing how well it forecasts the 2021 inflation surge. In contrast to standard benchmarks, which mechanically project a reversion to the mean, BISTRO correctly anticipates the persistence of the inflation wave. This highlights its ability to adapt to unfamiliar patterns in the data. Thus, BISTRO holds promise for producing reliable baseline forecasts and for scenario analysis.

JEL classification: C32, C45, C55, C87

Predictions of macroeconomic aggregates are a key ingredient to economic policy, especially for monetary policymakers. Traditionally, forecasting has relied on time series models, where historical developments of key economic variables and their correlations are used to extrapolate the future path of the variable(s) of interest. In this traditional approach, each model is tailored to the specific task at hand, meaning it has to be constructed, estimated and validated for a single problem. When dealing with another task involving a different set of variables in another setting, another model designed specifically for the new problem needs to be built. To use an analogy, the modeller is like a carpenter who has access to a large box of diverse tools. A new task is tackled by figuring out how to use or modify tools that were created for another task.

The advent of large language models (LLMs) has popularised a very different approach to problem-solving. The latest generation of LLMs are like a Swiss army knife rather than a box of tools: they can tackle a wide range of problems, even ones that were unforeseen when the model was built and trained. In the industry jargon, LLMs are "zero-shot learners" or "foundational models".

[1] The views expressed in this publication are those of the authors and not necessarily those of the BIS or its member central banks. We also thank Jon Frost, Gaston Gelos, Denis Gorea, Benoit Mojon, Matthias Rottner, Phurichai Rungcharoenkitkul, Dan Rees, Andreas Schrimpf, Frank Smets and Sonya Zhu for helpful comments and suggestions.

Key takeaways

- *We introduce the BIS Time-series Regression Oracle (BISTRO), a general purpose model for prediction of macroeconomic aggregates.*
- *BISTRO works as a "ChatGPT" for macroeconomic time series and provides accurate forecasts for key aggregates, and its flexibility enables a straightforward exploration of scenarios.*
- *To illustrate how BISTRO can help central bankers in their analysis, we showcase an application to inflation forecasting, in which we highlight its ability to uncover non-linear patterns.*

The ever-growing capabilities of LLMs have opened up the tantalising prospect that a similarly flexible approach would support time series forecasting. Rather than having to build a bespoke model for a particular task, the same foundational model could be deployed for a wide variety of different tasks.[2] The main advantage of such a model is that forecasting can break free of the inflexibility of traditional econometric approaches, in which specific modelling choices must be imposed from the start.

This article introduces such a flexible model. We name it the BIS Time-series Regression Oracle (BISTRO). BISTRO builds on the machinery underlying LLMs and applies it to the world of economic time series.

The article starts with a general introduction to the mathematical principles underlying LLMs and explores why they have broader applicability beyond the domain of text. Just as LLMs do well in guessing the next word within the broader context of the sentences that precede it, foundational time series models do well in guessing the next realisation of a macroeconomic time series within the broader context of what else has been happening in the economy. The only difference is that they operate on macroeconomic variables rather than words. But mathematically, the task is identical.

We then demonstrate how BISTRO can assist economists in their forecasting and scenario analyses. For example, a researcher can produce a generic baseline forecast for, say, inflation and then evaluate how conditioning on different explanatory variables (and different assumptions for their evolution) modifies the baseline. Hence, BISTRO constitutes a low-cost and easy-to-use forecasting tool that performs well compared with traditional econometric benchmarks. In particular, we show how BISTRO anticipates the persistence and pervasiveness of the inflation surge in 2021.

BISTRO comes with detailed instructions and pre-compiled scripts on how to operate it, available at https://github.com/bis-med-it/bistro. To facilitate replication and practical use, these scripts can be run in Google Colab, allowing users to upload their own data set and generate baseline and conditional forecasts with BISTRO through a guided workflow.

[2] Early proposals of automatic and flexible modelling for time series (Hendry and Krolzig (2001)) hinged on a necessarily task-specific model selection step.

From word prediction to foundational language models

Since the emergence of large language models (LLMs), economists have increasingly used them for text-based tasks such as drafting, coding and literature reviews.[3] Artificial intelligence (AI) tools are now routinely applied across many areas of economic analysis and research.[4]

LLMs are valuable because they can handle a wide range of natural language processing (NLP) tasks without any modification to their structure or parameters. Models such as the Generative Pretrained Transformer (GPT) can translate text, summarise documents and write code – even though they were trained for a single and very specific task: predicting the next word in a sequence.[5] This ability to perform tasks for which the model received no explicit training is known as *zero-shot learning* and is what makes LLMs *foundational models*: they are general purpose tools applicable to a broad set of tasks, possibly unforeseen when the model was created.

The zero-shot capability of transformers is rooted in how they represent words. In any quantitative model, words must first be converted into a vector of numbers, a process known as *embedding*. A simple approach would assign each word a fixed vector, regardless of context. But consider the word "bond": in isolation, it could refer to a financial instrument, a family connection or a fictional spy (BIS (2024)). A fixed, standalone vector cannot capture this range of use. Transformers address this limitation through *contextualised embeddings*. Rather than assigning a static vector to each word, the transformer reads the full surrounding text and constructs a vector that reflects the word's specific meaning in that specific context. The word "bond" would thus receive a different vector representation in a sentence about sovereign debt than in one about family relationships or spy fiction. This is achieved through the *attention mechanism*, a core component of the transformer architecture.[6] Attention allows each word to attend to every other word in the input, weighting their contributions according to relevance. The result is a set of embeddings that encode not just the identity of each word, but its role and meaning within the input sequence.

As the model generates a different internal representation for every new input, it effectively adapts to the questions it is prompted with. This adaptation to the context is at the root of its flexibility. Within the given model, the context-dependent embeddings allow it to behave differently for each prompt. The term "foundational model" refers to such adaptability (a Swiss army knife as opposed to a box of different tools). It is this adaptability that distinguishes transformer-based LLMs from the older generation of expert systems and machine learning models.

[3] Korinek (2023) surveys how generative AI can support economic research, while Korinek (2025) explores specific applications, including the replication of existing results.

[4] See for example Aldasoro et al (2025), Aquilina et al (2025); Cao et al (2024); Gambacorta et al (2024); Gorodnichenko et al (2023); Horton (2023); Kwon, Park, Perez-Cruz and Rungcharoenkitkul (2024); Kwon, Park, Rungcharoenkitkul and Smets (2025); Ludwig and Mullainathan (2024); Siano (2025); and Zarifhonarvar (2026).

[5] The transformer architecture is a deep neural network originally developed for language translation. For further details, see Vaswani et al (2017). GPT was introduced by Radford and Narasimhan (2018) and further evaluated in Radford et al (2019) and Brown et al (2020).

[6] Vaswani et al (2017) introduced the attention mechanism for transformers, and Alammar (2018) subsequently elucidated its inner workings; see the Annex for further details.

From bespoke time series models to a general purpose forecasting tool

Traditional prediction and forecasting methods, whether rooted in econometrics or machine learning, share a common feature: each task calls for its own model. The researcher must first specify the model's functional form f and then estimate its parameters – typically a low-dimensional vector – using historical data. Then the forecasting equation from these traditional models can be written as:

$$\hat{y} = f(\boldsymbol{x}; \hat{\theta}),$$

where y is the target variable, and $\boldsymbol{x}$ and $\hat{\theta}$ are, respectively, the data and the vector of estimated parameters on which the model depends, ie $\hat{\theta} = \theta(\boldsymbol{x})$. The model functional form $f(.;.)$ – be it a simple autoregressive model with one lag, such as an AR(1), a space-state model or even a random forest or a neural network – defines how the data included in the model specification are processed to produce parameter estimates and forecasts.

By contrast, a foundational model, such as BISTRO, takes a different approach. First, no functional form f is imposed on the data or the estimates θ derived from the data. Instead, it works with a very general function derived from the transformer model (see the Annex). Importantly, the foundational model takes as input not only the underlying data series $\boldsymbol{x}$ for the problem at hand but also parameters derived from other (potentially massive) data used in its training.

In the case of LLMs that work with text, their training set includes a substantial portion of text on the internet. Thus, when the LLM operates on a particular task involving a document, it employs data from outside for the specific task that it is asked to perform. In the case of BISTRO, the training set includes the large macroeconomic data set maintained by the BIS (see Box A). In this way, the predictive power of a foundational model derives from the training set of data that goes well beyond the data involved in a particular task.

More formally, we can contrast BISTRO with the traditional time series model by expressing the prediction $\hat{y}$ as:

$$\hat{y} = g(\boldsymbol{x}; \widehat{\boldsymbol{W}}).$$

Here, the function g is the general predictive function derived from the transformer model (see the Annex for details). This takes as the inputs not only the data $\boldsymbol{x}$ for the specific problem at hand but also the vector $\widehat{\boldsymbol{W}} = \boldsymbol{W}(\boldsymbol{Z})$, which depends on a potentially vast set of other data $\boldsymbol{Z}$ that was used to train the model.

Note that both traditional econometric models and foundation models rely on the same information set, $\boldsymbol{x}$, to predict future outcomes. The key difference lies in how they use this information. An econometric model operates in two steps: it first estimates the model parameters using $\boldsymbol{x}$ and then generates predictions conditional on the estimates $\hat{\theta}$. A foundational model, by contrast, operates in a single step: it maps $\boldsymbol{x}$ directly to predictions, relying entirely on weights, $\widehat{\boldsymbol{W}}$, that were estimated during pre-training and remain fixed at inference time.

The practical payoff for forecasters is considerable. With traditional methods, each change in the forecasting setup (eg adding a variable, lengthening the available sample, switching from a univariate to a multivariate model) calls for a new round of model-building and estimation. A foundational model subsumes all these cases. When provided with a single variable, it produces a non-linear autoregressive forecast

in which the order is adapted automatically by the attention mechanism that decides which lags to include. If explanatory variables are included in the input, the model refines its predictions to reflect their influence – all without any change to the model's structure or parameters. Because the internal weights adapt to the data prompted, the model also naturally accommodates time-varying relationships.

Introducing the BIS Time-series Regression Oracle

Our BISTRO builds on the MOIRAI (Woo et al (2024)) architecture. MOIRAI is a prominent example of a foundational model for time series forecasting, adapting the transformer architecture from NLP to numerical sequences.[7] Several features distinguish time series from text and require specific design choices. Words belong to a finite vocabulary; time series observations are realisations of continuous variables. Text follows a single reading order; time series can be sampled at different frequencies. Even more importantly, time series models may need to account for explanatory variables whose influence may change over time. MOIRAI addresses these challenges by grouping consecutive observations into fixed-length blocks (*patches*). A patch size of eight, for instance, means the model processes and predicts eight observations simultaneously. MOIRAI was trained on over 27 billion data points drawn from diverse domains (including energy, healthcare and economics/finance) with sampling intervals spanning one second to one year. It can incorporate up to 128 covariates, process variable patch sizes and handle multivariate time series.

While MOIRAI is a powerful general purpose tool, its broad training base means it is not optimised for any single application. As we will show below in the forecasting exercises, its predictive accuracy for macroeconomic variables is poor, particularly at longer time horizons or with covariates. We therefore constrained and fine-tuned MOIRAI to focus on macroeconomic time series and boost its usefulness for central banks, hence creating BISTRO.[8]

We made two main adjustments to build BISTRO. First, the model was limited to a daily frequency and a patch size of 32 days. This allows BISTRO to handle both unconditional forecasts – in which all variables are known only up to a given date – and conditional forecasts, in which some series are set to follow certain future paths. Second, the model was fine-tuned on the BIS macroeconomic data set (see Box A). All data were converted to a daily frequency by carrying forward the last known value until a new release or revision became available, ensuring that no future information entered the training process. The specific training procedure and model parameters are described in a companion working paper (Koyuncu et al (2026)).

BISTRO is therefore a foundational model tailored to macroeconomic time series typically observed at monthly or quarterly frequencies. It should not be regarded as the definitive answer to any forecasting or nowcasting exercise, but rather as a low-cost and easy-to-use tool that can quickly provide reasonable solutions to many

[7] Chronos 2.0 (Ansari et al 2025) is another foundational model comparable to MOIRAI. Carriero at al (2024) compare the performance of several univariate off-the-shelf foundational models for macroeconomic forecasting.

[8] Fine-tuning a foundational time series model is analogous to the process that turns a general purpose language model into a practical assistant, such as ChatGPT, Claude or Gemini. In that context, fine-tuning adjusts the model's parameters so that it responds conversationally, filters out unreliable information prevalent in its original training data and adopts a style suited to its users.

Box A

Details of the data set

The full internal data set is made up of 4,925 time series across 63 economies, observed from 1970 through 2024. The training and evaluation window is set at 1984–2024, which is intended to improve coverage consistency. The data set mirrors the real-world information set used by central bank economists and forecasters: indicators arrive at different frequencies and with publication lags. So the data set includes series with varying periodicity and release lags.

The series are arranged into macro blocks to ensure broad coverage of macro-financial conditions. The biggest categories are exchange rates (1,699 series), prices and earnings (1,130), domestic interest rates and bond yields (731) and national accounts (554). Smaller macro blocks broaden coverage, reaching real economy indicators as well as monetary and balance sheet indicators. The topics they cover include labour market data (325), monetary aggregates (235), output indicators (153), demand indicators (83) and items from central bank balance sheets and income statements (15).

The data set has a mixed-frequency structure, and it supports analysis across different time horizons. Of the total 4,925 series, 2,151 are monthly, 1,314 are quarterly, 861 are annual, 536 are daily, 61 are weekly, and two are semi-annual. Country coverage is broad, with dense coverage in several major economies and regions, thereby supporting cross-country comparisons. The series are spread across other advanced economies (1,020), Asia (368), Latin America (318), the G3 major advanced economies (309), emerging Europe (291) and Africa and the Middle East (206).

Standard transformations are used to make the data comparable across variables and to make them suitable for real-time forecasting. During training, BISTRO uses levels, differences, as well as logs and log differences, but only as needed for the indicator. The use of a common timeline, along with publication lags, is intended to ensure forecasts rely only on information available at the time. Forecasts keep the most recent observation for each series and then extend that forward to the next release date to align the series. The approach respects publication lags.

forecasting and nowcasting problems. It can also serve as a benchmark: if a purpose-built econometric model cannot improve on BISTRO's predictions, it might call for additional model exploration. And, like any foundational model, BISTRO can be further fine-tuned, for instance to focus on a specific country or to address a particular forecasting problem.

How to operate BISTRO

BISTRO was designed as a ready-to-use and off-the-shelf tool for macroeconomic forecasting. This section describes the workflow, which comprises three stages: data preparation, parameter modelling, and forecasting and evaluation (Graph 1). A Google Colab notebook supports replication, with sample scripts and step-by-step code for the entire workflow, including rolling-window forecast generation.

In the data preparation step, the user aligns the target variable (eg inflation or unemployment) with selected covariates (eg oil prices or exchange rates) along the time dimension. Indicators released at different frequencies are converted to daily by carrying forward the last known value. This ensures that the data set contains only information available as of each date. For example, second quarter GDP appears in

Workflow for BISTRO

Graph 1

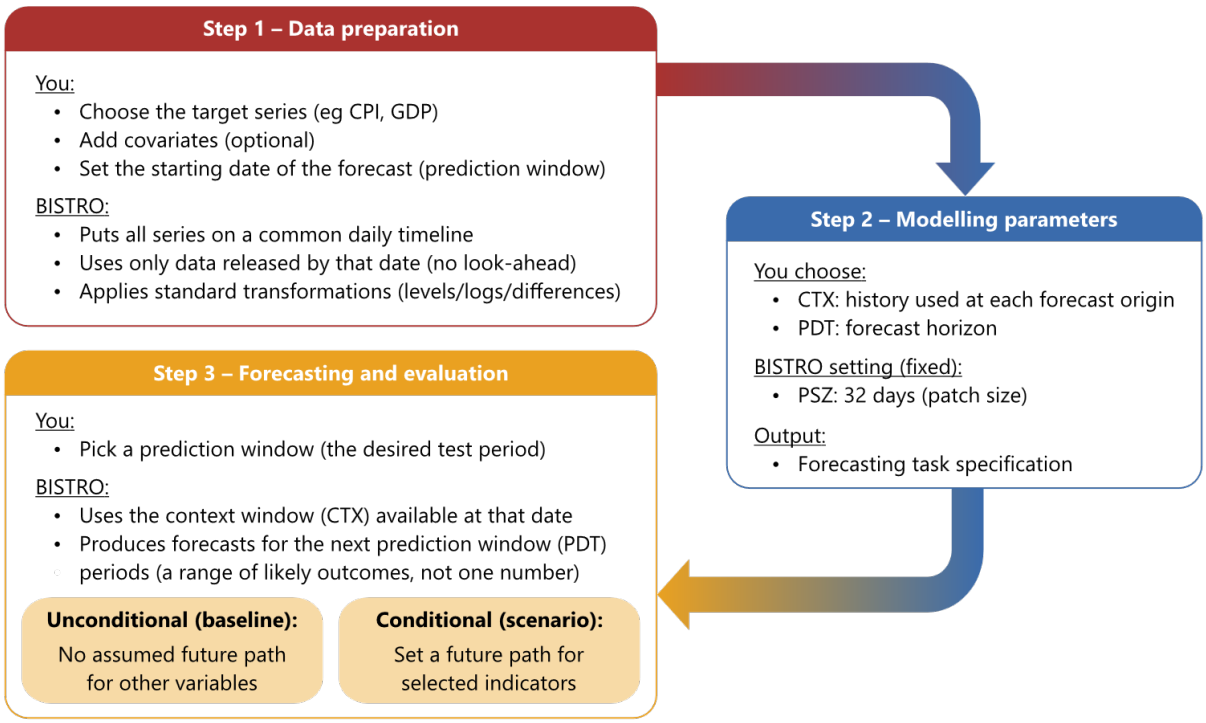

CTX = context window length; PDT = prediction window length; PSZ = patch size (fixed to 32).

The workflow has three steps. Step 1 builds the real-time data set for each forecast origin (forecast date): all series are put on a common daily timeline, and BISTRO uses only data released by that date (no look-ahead), carrying forward the last available value until the next release. Step 2 defines the task by choosing context length (CTX) and prediction length (PDT). Step 3 runs a rolling-origin back test over the evaluation window, produces probabilistic forecasts and checks accuracy using out-of-sample errors (forecast errors) and standard measures (eg root mean square error, or RMSE). Forecasts can be unconditional (baseline) or conditional (scenario) based on an assumed path for selected indicators.

Source: Authors' elaboration.

the data set only on its release date in the third quarter.[9] The result is a *(pseudo) real-time information set* that mirrors the operational environment faced by central bank modellers and prevents forecasts from using data that were not yet available.

The aligned data set is then passed to the model interface, which automatically handles the required preprocessing. It applies standard data transformations (levels, logs, differences), imputes missing values due to publication lags and determines the number of past observations used for each prediction, based on the user-defined context window length (CTX). For instance, a context window of 100 patches corresponds to roughly 8.75 years of historical data. Users do not need to manually transform or reformat the data.

A simple example illustrates how running the model mirrors operational forecasting. Suppose the goal is to forecast inflation from 2023 through end-2025, estimating the next 12 months of inflation (12 patches of 32 days) at each forecast date, using 20 years of history (228 patches). BISTRO is first prompted with observations from 1 January 2003 through 31 December 2022 and then predicts inflation for the next 12 months. It then shifts the data window forward by one patch,

[9] For simplicity, we abstract from data revisions. While revisions are an important issue for GDP, most other macroeconomic time series are revised only sporadically and by small amounts.

Forecasting modes in BISTRO: baseline and scenario[1] Graph 2

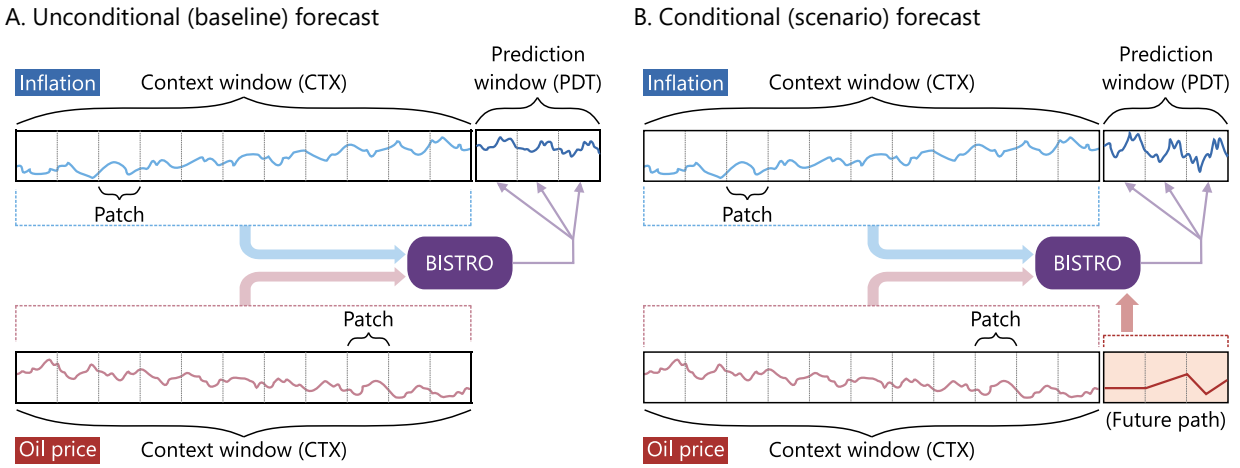

[1] In panel A, BISTRO produces a baseline inflation forecast using only information available at the forecast date. In panel B, the inflation forecast is updated by imposing an assumed path for oil prices over the forecast horizon. The difference between the two forecasts reflects historical co-movements in the data, not causal effects.

Source: Authors' elaboration.

using data from 1 February 2003 through 31 January 2023 to forecast the next 12 months. This process repeats until the end of the sample.

BISTRO automates the entire estimation process and calculates prediction errors at each step. It benchmarks the results against an AR(1) model, enabling a quick assessment of forecast quality relative to a common benchmark. In the example, the user supplies consumer price index (CPI) inflation data for the country or area of interest over the period 2003–25, sets the CTX to 228 patches and the prediction window length (PDT) to 12 patches. The AR model is trained on data through December 2022 and is updated monthly, while BISTRO uses all 228 context patches at each forecast origin without retraining.

To produce forecasts that also depend on covariates such as oil prices or exchange rates, the user adds these series and aligns them with the target variable. BISTRO then conditions its prediction on both the past values of the target and the explanatory variables (Graph 2.A). BISTRO can also produce conditional forecasts. To this end, the user fixes the future path of one or more covariates and feeds them to the model (Graph 2.B). For example, BISTRO can forecast 2026 inflation under alternative oil price trajectories. This allows researchers to explore different macroeconomic scenarios in a straightforward way.

Forecasting inflation with BISTRO

This section illustrates how BISTRO can be used to forecast inflation. More specifically, we compare the results obtained by prompting BISTRO with historical inflation data against those obtained from a standard time series benchmark – namely a naive AR(1)

– as well as against those produced by the baseline MOIRAI.[10] We first examine the accuracy of BISTRO's unconditional forecasts – ie simple out-of-sample predictions that do not depend on covariates. Then we showcase how forecasts are affected by the inclusion of explanatory variables (covariates) and how they can be tweaked to reflect assumptions on the future evolution of said covariates. This is typically labelled a conditional forecasting exercise.

The unconditional forecasting exercise mimics a task of paramount relevance for central banks – predicting monthly inflation at short and medium horizons. To have BISTRO (and MOIRAI) generate unconditional forecasts, we prompt it with the past values of inflation, starting 240 months before the beginning of the forecast. Note that BISTRO relies on the same amount of information as provided for a benchmark AR(1) – that is the latest inflation readings. The AR(1) model uses 240 months of past inflation values to estimate its parameters and then applies exponentially decaying weights to the last observation to predict the next 12 months. In contrast, BISTRO directly predicts the next 12 months of inflation using the very same 240 months of data. Yet this does not require parameter estimation: BISTRO relies on pre-estimated parameters from its generic training process.

To illustrate BISTRO's forecasting capabilities, we first put it to test on the inflation wave of 2021–22.[11] Graph 3 shows US inflation forecasts made by BISTRO,

US Inflation forecasts throughout the 2021 surge[1]

In per cent

Graph 3

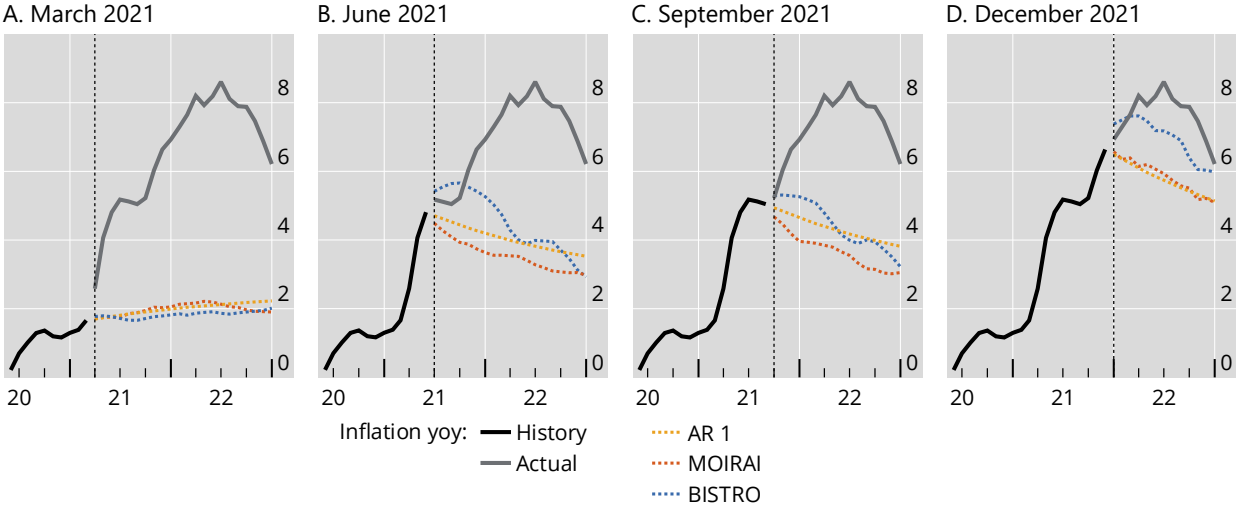

[1] The graph presents a comparison of inflation forecasts extending until the end of 2022, as produced by BISTRO (blue), MOIRAI (orange) and AR(1) (yellow). The observed inflation is denoted in black, while the unobserved (by the models) realised inflation is indicated in grey. Each panel displays inflation estimates commencing from a specific date during the period of the US inflation surge in 2021 (March, June, September and December). Notably, the panels for June and December in particular highlight BISTRO's ability to accurately forecast the inflation peaks.

Sources: BIS Consumer prices; authors' calculations.

[10] Despite their simplicity, autoregressive models are a tough benchmark to beat when forecasting macroeconomic aggregates. See for example Faust and Wright (2013) and Hall et al (2023).

[11] Admittedly, this period was also included in the training sample. Yet note that in the training of the model, the inflation time series was never presented in isolation or explicitly labelled as "inflation", and it was expressed in month-on-month rather than year-on-year terms. Furthermore, the only information available to the model consisted of the numerical values within the time series. Had the

MOIRAI and a benchmark AR(1) model, starting at different points of the inflation surge. Each panel shows multi-step-ahead forecasts, in which predictions are produced dynamically over the desired forecast horizon and hence do not rely on inflation readings that were not yet available. In March 2021, when inflation had not yet surfaced, all models provide a similar forecast of inflation converging towards the 2% target (Graph 3.A). Yet as the first signs of the inflation surge appear, in June 2021, the model projections diverge (Graph 3.B). The AR(1) benchmark mechanically decays towards the historical average, while the MOIRAI model hints at an even faster decline. By contrast, BISTRO correctly anticipates a more persistent wave of inflation. In September 2021 inflation had apparently plateaued, and while the AR(1) and MOIRAI stick to a declining trajectory, BISTRO hints at a shallower decline (Graph 3.C). In December, a second wave of inflation had already taken hold, and BISTRO foresees rising and persistently high inflation, close to the actual outcome (Graph 3.D).[12]

To be sure, the 2021–22 wave of inflation is a very specific episode which simple linear time series models have a hard time processing. The exercise above was indeed meant to illustrate how BISTRO can identify and deal with possible structural changes and regime shifts. Yet BISTRO performs competitively and consistently over different time periods, not necessarily featuring structural changes and extreme circumstances. To establish this, we run a full-fledged forecasting comparison based on four distinct evaluation windows: 1995, 2005, 2015 and 2023–24, which were not included in the BISTRO's training sample. These windows reflect markedly different macroeconomic regimes, ranging from the pre- and post-Great Financial Crisis period to the recent episode of elevated inflation.

This forecasting setup allows us to assess model performance across phases characterised by different levels of volatility, sources of price pressures and potential structural changes. Within each evaluation window, we focus on forecasts of the year-on-year (yoy) growth in the CPI at one-, three-, six- and 12-month-ahead horizons (PDT = 1, 3, 6, 12). Forecasts are generated sequentially over the test window as it progresses forward in time.[13] The analysis is conducted for all countries for which sufficient data are available. Forecasts are separate for each country, but for ease of presentation, results are aggregated into four groups: the United States (US), euro area, other advanced economies (other AEs) and emerging market economies (EMEs).[14]

The results show that BISTRO provides reliable and accurate forecasts compared with the benchmark. We measure forecast accuracy using relative root mean squared forecast errors (R-RMSFE) with respect to the AR(1) benchmark, averaged across countries and testing windows within each group (Table 1). So a value below one indicates that BISTRO (or MOIRAI) outperforms the benchmark. On average, BISTRO improves upon the benchmark in most cases, with particularly strong gains at the longer horizons. For the United States, BISTRO outperforms both AR(1) and MOIRAI for all horizons except the one-month-ahead, with its relative performance improving

model memorised the time series exactly, it would have perfectly replicated the forthcoming surge in inflation. That said, later in the article we also run a full-fledged forecasting comparison using forecasting windows that were deliberately excluded from the training sample.

[12] Faria-e-Castro and Leibovici (2024) employ a text-based LLM to predict the 2021 US inflation surge, achieving results comparable to AR(1) predictions.

[13] To ensure comparability across models, we apply the same principle to the AR(1) benchmark, whose coefficients are not iteratively re-estimated throughout the test period.

[14] For a detailed breakdown of the results, see Koyuncu et al (2026).

Forecast accuracy of BISTRO across regions and evaluation windows

1995, 2005, 2015 and 2023–24 (R-RMSFE; PDT = 1, 3, 6, 12) Table 1

Region	Horizon 1 month		Horizon 3 months		Horizon 6 months		Horizon 12 months	
	BISTRO	MOIRAI	BISTRO	MOIRAI	BISTRO	MOIRAI	BISTRO	MOIRAI
US	1.024	1.136	0.946	1.241	0.804	1.335	0.481	1.349
Euro area	1.065	0.987	0.927	0.932	0.890	0.956	0.898	1.689
Other AEs	1.157	1.053	1.052	1.108	0.972	1.220	0.813	1.097
EMEs	1.009	0.874	0.876	0.844	0.895	0.854	0.969	0.850

For each forecast horizon (PDT), we calculate the root mean squared forecast errors (RMSFE) for BISTRO, MOIRAI and the AR(1) benchmark. We then compute the ratio of RMSFE values of BISTRO and MOIRAI relative to those of the AR(1) benchmark and report their average across evaluation windows and countries within each region.

AEs = advanced economies: Australia, Canada, Czechia, Denmark, Hong Kong, Israel, Japan, Korea, New Zealand, Singapore, Sweden, Switzerland, United Kingdom.

EMEs = emerging market economies: Algeria, Argentina, Brazil, Chile, China, Colombia, Hungary, India, Kuwait, Malaysia, Mexico, Morocco, Peru, Philippines, Poland, Romania, Russia, Saudi Arabia, South Africa, Thailand, Türkiye, Vietnam.

Source: Authors' calculations.

as the forecasting horizon lengthens. This is similar for the euro area, where both general purpose models perform similarly on average and improve with the length of the forecasting horizon. Results for other advanced economies are mixed, with average performance closer to the AR(1) benchmark.[15] Similarly, for EMEs, BISTRO delivers systematic improvements at the medium-term horizon.

Constructing scenarios with BISTRO

Having established the solid performance of our model in unconditional forecasting, we turn to scenarios, that is, conditional modelling exercises. One of the main advantages of the flexibility of our model lies in its ability to naturally handle relevant covariates that one may want to condition the forecast on. In standard time series models, conditioning is possible only on the set of variables that were considered and explicitly included in the model's specification and estimation. We can instead prompt BISTRO with trajectories of additional macroeconomic variables (not necessarily featured in the training sample) and then generate projections that are conditional on their realisations. The way the model handles this is by relying on the historical patterns that emerge from the training data set and hence shape the embeddings.

To illustrate this, we consider conditional forecasts of inflation based on different possible future trajectories of the oil price. Starting from its realised path, which serves as a baseline, we introduce counterfactual oil price scenarios in which the oil price would be 3% lower and 3% and 9% higher than it actually was. In general, conditioning on oil prices produces systematically higher or lower oil price paths relative to the baseline (Graph 4.A). But even more importantly, the model unveils a substantial degree of non-linearity in the reaction of inflation to oil prices. First, a decline in the price of oil has smaller medium-run effects on inflation than an increase

[15] This is partly because BISTRO makes quite large errors compared with the AR(1) benchmark for the United Kingdom and New Zealand in the 2023–24 test window.

US inflation counterfactual forecasts under oil price scenarios[1]

In percentage points

Graph 4

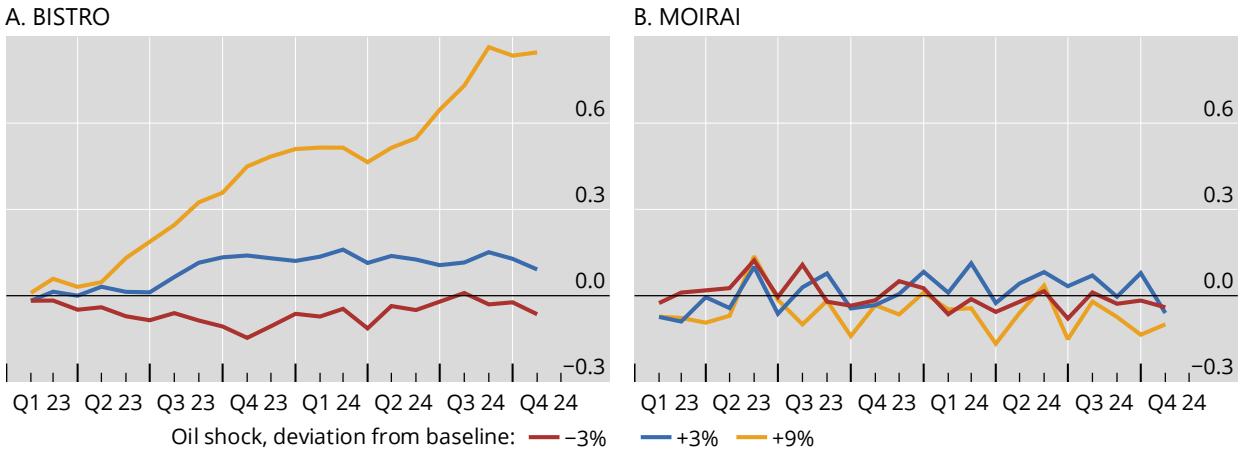

[1] The figure compares counterfactual year-on-year inflation projections from BISTRO (panel A) and MOIRAI (panel B) under alternative oil price paths (baseline and shocks of −3%, 3% and 9%).

Sources: Bloomberg; BIS Consumer prices; author's calculations.

by the same amount. Second, larger increases in the price of oil have a more than proportional effect on inflation. A clear bearing on the relationship between oil prices and inflation is one of the advantages of our model over MOIRAI. The latter, trained on a wider set of non-macroeconomic time series, fails to highlight any meaningful connection between oil prices and inflation (Graph 4.B).

Monetary policy as viewed through the prism of BISTRO

While our model is a device to uncover (potentially non-linear and complex) correlations in the data, it relies entirely on the historical patterns that emerged during the training. Hence, it cannot provide insights into the effects on the observed time series of unobservable macroeconomic shocks, nor can it shed light on how structural forces with an economic interpretation contribute to its forecasts. So, for example, if one generated counterfactual scenarios based on different trajectories of monetary policy, one would see that higher policy rates are associated with higher inflation. This is because the model has no understanding of a policy rule or of how a "tighter" policy transmits to the real economy. Hence, conditioning on a policy rate trajectory will yield only a correlation based on historical patterns.

However, this suggests that the model can be used to gain insights on historical response functions, eg a Taylor rule for monetary policy. To do this, we turn to conditional forecasting exercises in which the model is provided with realised macroeconomic paths and asked to generate the implied policy rate projections based on such a macroeconomic outlook. More precisely, we condition the forecast for the US policy rate on the realised path of inflation and examine the implied trajectories of the policy rate obtained by providing different historical contexts. This design allows us to isolate how the historical information provided affects the model-embedded relationship between macroeconomic conditions and monetary policy.

The implied US policy rate forecast, conditioning on inflation[1]

In per cent

Graph 5

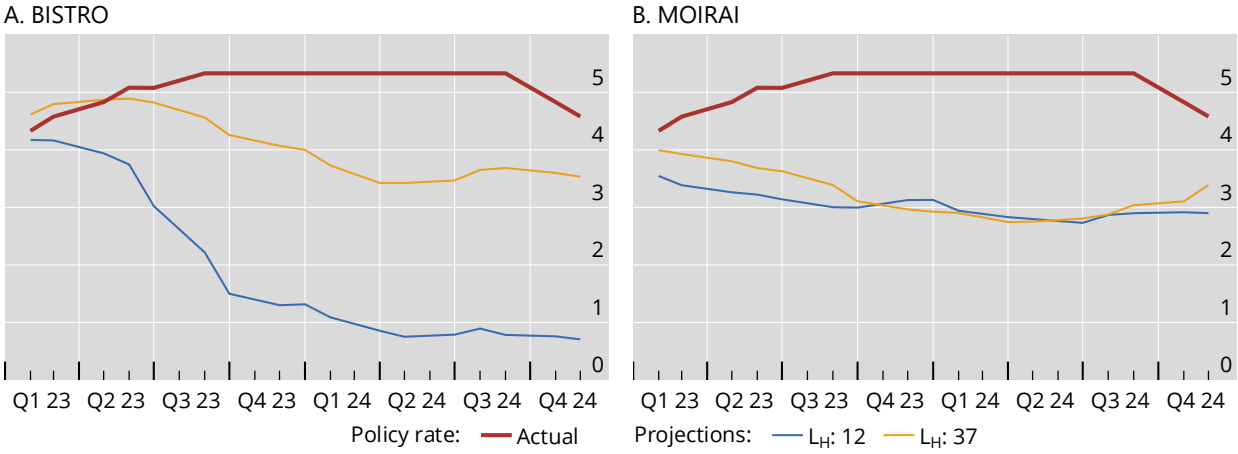

L_H = length of the historical sample (in years).

[1] The graph compares policy rate projections from BISTRO (panel A) with those from MOIRAI (panel B) using 12-year versus 37-year historical samples.

Sources: BIS Central bank policy rates; authors' calculations.

Results crucially depend on the amount of historical information we condition on (Graph 5). When conditioning only on the last 12 years of data in our sample (eg 2011–22), the model uses information based on a low-inflation regime in which monetary policy was mostly conducted through unconventional tools, due to the effective lower bound. Hence, the link between policy rates and inflation appears quite weak. As a result, the model fails to recognise that a policy tightening is necessary to tame inflation and instead hints at near-zero policy rates, implicitly assuming that inflation would self-correct without the need for policy intervention.

If one instead extends the context window to cover the entire training sample (1986–2022), the model gains exposure to different inflationary regimes and to a phase in which conventional policy played a key role keeping inflation in check. In this richer context, the model correctly infers that high inflation typically calls for a tightening of policy rates, leading to a somewhat sharper rise in the implied interest rate path that slightly exceeds the realised pace of tightening in the very beginning. By contrast, the projections generated by MOIRAI are largely insensitive to the length of the conditioning history, yielding similar policy rate paths across the two settings. This contrast suggests that BISTRO makes more effective use of historical information when capturing the correlations between inflation and monetary policy.

Opening the black box: the role of the attention weights

The different performances of BISTRO and MOIRAI in dealing with covariates and conditional scenarios underscore the key role played by the attention weights. While BISTRO's attention weights effectively integrate covariate data to generate meaningful scenarios, MOIRAI fails to assimilate this information, resulting in less informative projections.

To explain why this is so, let us first focus on MOIRAI. As illustrated in Graph 4.B, MOIRAI's predictions remain unchanged across the four distinct oil price scenarios. This suggests that the attention weights assigned to the oil price time series are effectively zero, therefore the model ends up relying entirely on historical inflation data. Consequently, MOIRAI is unable to infer the relationship between inflation and oil prices from past observations. Furthermore, Graph 5.B demonstrates that the predicted policy rate differs only marginally when contrasting short and long historical contexts. This implies that MOIRAI predominantly relies on the most recent readings, largely disregarding earlier values of both the policy rate and inflation.

In contrast, BISTRO, having been specifically trained on macroeconomic time series data, demonstrates an ability to assign attention weights that capture relevant historical information. In Graph 4.A, the model's non-linear attention weights successfully reflect the complex, non-linear impact of oil prices on inflation. Additionally, as shown in Graph 5.A, BISTRO leverages longer time horizons by assigning greater attention to observations from the more distant past, thereby enhancing its policy rate projections. This adaptability of the transformer model, whereby it can refine its estimates by adjusting attention weights, rather than requiring retraining, underscores its strength as a general purpose forecasting tool. Nevertheless, this also highlights the necessity for the model to be exposed to time series pertinent to the specific forecasting task.

Concluding remarks

The key advantage of BISTRO compared with alternatives lies in its flexibility. It can accommodate prediction of a multitude of time series in unconditional as well as conditional terms. Importantly, it is also able to unveil and deal with non-linearities. While we provided a few examples, BISTRO lends itself to plenty of different empirical exercises, ranging from pure forecasting to scenario analysis.

References

Alammar, J (2018): "The illustrated transformer", June, https://jalammar.github.io/illustrated-transformer/.

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

Ansari, A F, O Shchur, J Küken, A Auer, B Han, P Mercado, S S Rangapuram, H Shen, L Stella, X Zhang, M Goswami, S Kapoor, D Maddix, P Guerron, T Hu, J Yin, N Erickson, P M Desai, H Wang, H Rangwala, G Karypis, Y Wang and M Bohlke-Schneider (2025): "Chronos-2: from univariate to universal forecasting", arXiv:2510.15821, October.

Aquilina, M, D Araujo, G Gelos, T Park and F Perez-Cruz (2025): "Harnessing artificial intelligence for monitoring financial markets", *BIS Working Papers*, no 1291, September.

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

Brown, T, B Mann, N Ryder, M Subbiah, J Kaplan, P Dhariwal, A Neelakantan, P Shyam, G Sastry, A Askell, S Agarwal, A Herbert-Voss, G Krueger, T Henighan, R Child, A Ramesh, A Ziegler, J Wu, C Winter, C Hesse, M Chen, E Sigler, M Litwin, B Gray, B Chess, J Clark, C Berner, S McCandlish, A Radford, I Sutskever and D Amodei (2020): "Language models are few-shot learners", *Advances in Neural Information Processing Systems*, vol 33.

Cao, S, W Jiang, J L Wang and B Yang (2024): "From man vs. machine to man + machine: the art and AI of stock analyses", *Journal of Financial Economics*, vol 160, no 103910.

Carriero, A, D Pettenuzzo and S Shekhar (2024): "Macroeconomic forecasting with large language models", arxiv.org 2407.00890, July.

Faria-e-Castro, M and F Leibovici (2024): "Artificial intelligence and inflation forecasts", *Federal Reserve Bank of St. Louis Review*, vol 106, no 12.

Faust, J and J H Wright (2013): "Forecasting inflation", *Handbook of Economic Forecasting*, vol 2, part A, chapter 1.

Gambacorta, L, B Kwon, T Park, P Patelli and S Zhu (2024): "CB-LMs: language models for central banking", *BIS Working Papers*, no 1215, October.

Gorodnichenko, Y, T Pham and O Talavera (2023): "The voice of monetary policy", *American Economic Review*, vol 113, no 2.

Hall, S, G Tavlas and Y Wang (2023): "Forecasting inflation: the use of dynamic factor analysis and nonlinear combinations", *Journal of Forecasting*, vol 42, no 3.

Hendry, D and H-M Krolzig (2001): *Automatic econometric model selection using PcGets 1.0*, Timberlake Consultants Press.

Horton, J (2023): "Large language models as simulated economic agents: what can we learn from Homo Silicus?", *NBER Working Papers*, no 31122, April.

Korinek, A (2023): "Generative AI for economic research: use cases and implications for economists", *Journal of Economic Literature*, vol 61, no 4.

——— (2025): "AI agents for economic research", *NBER Working Papers*, no 34202, September.

Koyuncu, B, B Kwon, M Lombardi, F Perez-Cruz and H S Shin (2026): "Introducing BISTRO: a foundational model for unconditional and conditional forecasting of macroeconomic time series", *BIS Working Papers*, forthcoming.

Kwon, B, T Park, F Perez-Cruz and P Rungcharoenkitkul (2024): "Large language models: a primer for economists", *BIS Quarterly Review*, December.

Kwon, B, T Park, P Rungcharoenkitkul and F Smets (2025): "Parsing the pulse: decomposing macroeconomic sentiment with LLMs", *BIS Working Papers*, no 1294, October.

Ludwig, J and S Mullainathan (2024): "Machine learning as a tool for hypothesis generation", *Quarterly Journal of Economics*, vol 139, no 2.

Radford, A and K Narasimhan (2018): "Improving language understanding by generative pre-training", OpenAI, mimeo.

Radford, A, J Wu, R Child, D Luan, D Amodei and I Sutskever (2019): "Language models are unsupervised multitask learners", OpenAI, technical report.

Siano, F (2025): "The news in earnings announcement disclosures: capturing word context using LLM methods", *Management Science*, vol 71, no 11.

Vaswani, A, N Shazeer, N Parmar, J Uszkoreit, L Jones, A N Gomez, Ł Kaiser and I Polosukhin (2017): "Attention is all you need", *Advances in Neural Information Processing Systems*, vol 30.

Woo, G, C Liu, A Kumar, C Xiong, S Savarese and D Sahoo (2024): "Unified training of universal time series forecasting transformers", *ICML'24: proceedings from the 41st International Conference on Machine Learning*, article no 2178.

Zarifhonarvar, A (2026): "Generating inflation expectations with large language models", *Journal of Monetary Economics*, vol 157(C).

Mikael Juselius
mikael.juselius@bis.org

Dora Xia
dora.xia@bis.org

Monetary responses to external shocks in emerging market economies: the role of financial vulnerabilities[1]

We examine how reliance on external funding and foreign exchange (FX) market depth shape monetary policy responses to external shocks in emerging market economies (EMEs) with floating exchange rate arrangements. Using a panel of EMEs, we document that central banks in economies with both high external foreign currency debt and shallow FX markets move their policy rates in step with US monetary policy surprises. Notably, these countries show smaller exchange rate moves compared with other countries, consistent with the policy actions dampening FX fluctuations. Other economies, including those with larger external local currency debt, do not display these patterns. We also find that EMEs with shallow FX markets tend to deploy FX intervention following US monetary policy surprises.

JEL classification: E58, F41, F42

Reliance on external financing can complicate monetary policy in emerging market economies (EMEs) via its effects on exchange rates. This is because the standard textbook prescription – that floating exchange rates can absorb the external spillovers, leaving monetary policy free to respond to output and inflation – does not necessarily hold when underlying financial market frictions amplify the destabilising macroeconomic effects of external shocks (Basu et al (2025)). In this case, the appropriate policy response can differ from the textbook one and may be complemented by other instruments that are tailored to address specific exposures and vulnerabilities (BIS (2019, 2023); IMF (2020)).[2]

Two key financial frictions can weaken or even reverse standard monetary transmission in EMEs: shallow foreign exchange (FX) markets and limits on external borrowing. For countries reliant on external financing, shallow FX markets make it difficult for EME borrowers and global lenders to effectively hedge currency exposures, resulting in currency mismatches that directly affect balance sheets. In turn, balance sheets constrain borrowing capacity. Together, these two frictions create a feedback loop that amplifies the impact of exchange rate movements on

[1] The views expressed in this publication are those of the authors and not necessarily those of the BIS or its member central banks. We are grateful to Jimmy Shek for excellent research assistance. We also thank Julián Caballero, Paolo Cavallino, Gaston Gelos, Denis Gorea, Benoît Mojon, Daniel Rees, Andreas Schrimpf, Frank Smets, Ilhyock Shim, Hyun Song Shin, Vlad Sushko, Philip Woodridge and Christian Upper for helpful comments and suggestions.

[2] See also Brandão-Marques et al (2021) and Gelos and Sahay (2023).

Key takeaways

- *EME monetary policy responses to external shocks are shaped by reliance on external funding and FX hedging market depth.*
- *EMEs with higher foreign currency debt and less developed FX markets tend to move their interest rates in step with monetary surprises in the United States, whereas other countries do not.*
- *EMEs with higher foreign currency debt and shallower FX markets exhibit smaller FX moves in response to foreign monetary policy surprises, consistent with policy actions dampening FX fluctuations.*

financial conditions. A concrete example is when EME borrowers hold unhedged foreign currency liabilities. If the local currency depreciates, the domestic value of those liabilities rises, reducing borrowers' net worth and tightening credit limits. Similar channels operate on the supply of credit, via global banks (Bruno and Shin (2015)) and global portfolio investors (Carstens and Shin (2019)). When these channels are sufficiently strong, a reduction in the domestic policy rate can tighten financial conditions and even become contractionary, as it is accompanied by currency depreciation.

Although the mechanisms through which external shocks affect EMEs are well mapped out, far less is known about how EME policymakers respond in practice.[3] Such responses offer important clues about the constraints, priorities and trade-offs they face. Evidence to date suggests responses are varied and depend on the context. Mano and Sgherri (2024) examine how EME policy rates, FX interventions and exchange rates react to capital flow shocks. They find mixed policy responses across countries and present tentative results that choices reflect country-specific vulnerabilities. Mendoza-Fernández and Pelin (2025) study EME reactions to US monetary shocks and find that central banks move rates in the same direction as the shock.[4] They connect this reaction to import price pass-through to inflation.

In this article, we study how high exposure to foreign liabilities and shallow FX markets in EMEs shape the monetary policy response to external financing shocks. These aspects have not been explicitly incorporated in previous studies.[5] While we do not directly discuss trade channels, we take them into account in our estimations.

We document that both large external liabilities and shallow FX markets are crucial for understanding EME policy responses to external financial shocks. Economies exhibiting both vulnerabilities tend to move policy rates even more than one for one in the same direction as the US monetary policy surprise, whereas others do not. Moreover, economies with shallower FX markets additionally use FX interventions to provide liquidity to thin markets and to dampen excessive and

[3] By contrast, spillover effects on macro-financial outcomes are extensively documented; see Claessens et al (2016) and Kearns et al (2023), and references therein.

[4] This aligns with Calvo and Reinhart (2002) on "fear of floating", whereby interest rates are used to dampen exchange rate volatility. Related work estimates Taylor-type rules to test similar hypotheses, eg Montes and Ferreira (2020) and De Leo et al (2024).

[5] Our analysis is close to Mendoza-Fernández and Pelin (2025) in terms of methodological design, as we use US monetary shocks as a proxy for exogenous external financing shocks. In line with Mano and Sgherri (2024), we consider the mix between interest rates and FX interventions in the responses.

destabilising currency fluctuations, regardless of the level of their foreign debt exposures.

External liabilities and monetary policy transmission

In this section, we examine how the structure of foreign borrowing can complicate monetary policy in EMEs via the exchange rate. A key challenge is that when EME central banks adjust the policy rate in response to external shocks, the outcome can be the opposite of what standard textbook prescriptions would suggest, depending on which transmission channels dominate. For instance, lowering rates to support demand can instead lead to an exchange rate depreciation that tightens financial conditions and reduces demand, as discussed in detail below. In these circumstances, monetary policy may need to move rates in the same direction as the shock or rely on complementary tools, such as FX interventions and macroprudential policies, to support demand and stabilise financial conditions (BIS (2019)).

Two underlying market frictions are pivotal in reversing the traditional monetary policy transmission (Basu et al (2025)).[6] The first is shallow FX markets. External financing often creates currency mismatches, on either borrowers' balance sheets (eg an EME corporate borrowing in foreign currency) or lenders' balance sheets, such as foreign investment in local currency bonds. When FX markets are shallow, hedging instruments are limited and costly, making effective hedging impractical and leaving currency mismatches remaining on the balance sheets. The second friction is that EME borrowing is limited by the balance sheet strength of these borrowers and lenders. As their balance sheets deteriorate, their capacity to borrow or lend is curtailed. Together, these two frictions link financial conditions to the exchange rate. Specifically, depreciation weakens balance sheets, tightens financial conditions and reduces demand, while appreciation has the opposite effect. As EME monetary policy loosening commonly gives rise to EME currency depreciation, the subsequent tightening of financial conditions can counteract its usual expansionary effects.

The literature has identified several specific transmission channels depending on whose balance sheets are exposed to currency mismatches. While the two frictions above underpin all of them, they have different implications for our empirical setup. We discuss each in turn.

One channel arises when EME borrowers borrow in foreign currency – either because they cannot issue in local currency (the "original sin" of Eichengreen and Hausmann (1999)) or because they opt to borrow in foreign currency. This involves mismatches on their balance sheets, where foreign currency liabilities are paired with local currency assets.[7] If the currency mismatch cannot be effectively hedged, an exchange rate depreciation, say, would increase the domestic value of these debts, and consequently lower borrowers' net worth. This would lead to higher debt servicing burdens and tighter borrowing limits, both of which reduce demand. EMEs with large (unhedged) liabilities that are denominated in foreign currency should be more vulnerable to this channel.

[6] Checo et al (2024) find that domestic monetary policy shocks in EMEs are transmitted in the same way as in advanced economies. For instance, monetary loosening is expansionary.

[7] Céspedes et al (2004), Cook (2004), Gourinchas (2018) and Akıncı and Queralto (2024) study the theoretical underpinnings of this channel.

Deterioration of borrowers' balance sheets can also affect the supply of credit from global lenders. The reason is it increases exposures to credit risk for global banks which respond by reducing lending and raising spreads. This impact is known as the financial channel of the exchange rate (Bruno and Shin (2015)). Disentangling this channel requires data on foreign currency loans by global banks and associated credit spreads.

Yet another mechanism arises from unhedged positions of global portfolio investors who hold local currency assets while leaving currency risk unhedged. When EME currencies depreciate, there is an amplification of any losses on these assets in foreign currency terms. The investor loses both from the decline of asset values in local currency terms and the depreciation of the local currency. Such episodes are associated with portfolio outflows from EMEs and tightening of local financial conditions. This dynamic can create a self-reinforcing loop between deteriorating financial conditions and further EME currency depreciation. Carstens and Shin (2019) lay out the details of this channel, which they refer to as "original sin redux". EMEs with a large amount of local currency bonds held by foreigners should be more exposed to this channel. As before, hedging would in principle help mitigate it, but shallow FX markets and elevated hedging costs make this difficult in practice.

Some channels also reinforce the normal textbook monetary policy transmission via the exchange rate. For instance, some EMEs are increasingly investing in foreign currency-denominated assets. When these exposures are unhedged, exchange rate moves reinforce domestic monetary transmission: a policy rate cut that leads to a currency depreciation raises the local currency value of foreign assets, strengthening EME investors' balance sheets and expanding their domestic credit supply (Shin et al (2025); Juselius et al (2025)). This is original sin redux, but in reverse with EMEs and advanced economies switching net lender and borrower positions.[8]

In sum, the balance between different channels determines how changes in domestic policy rates are transmitted to the economy and therefore affects how central banks respond to external shocks. Generally, EMEs that are more exposed to unhedged foreign funding and where FX markets are less developed face greater challenges in countering the impact of external financing shocks on output and inflation. In contrast, transmission is more straightforward in EMEs with deeper FX markets and greater foreign assets. All these considerations have implications for our empirical setup, which we discuss next.

EME policy responses to foreign monetary shocks

We now turn to data to explore if actual EME monetary policy reactions depend on the channels described above. To this end, we have collected measures for: (i) the policy instruments; (ii) external financing shocks; (iii) exposures to currency mismatch from foreign funding and FX market depth; and (iv) controls and other outcome variables.

The policy instruments that we consider are the policy rate and FX interventions. The policy rate is available at a daily frequency. Because accurate FX intervention

[8] In addition, the exchange rate also affects trade and imported inflation, both of which matter for the inflation output trade-off facing EME central banks. The strength of transmission depends on whether exports and imports are priced in local currency or not (Gopinath et al (2020)).

measures are hard to obtain, we use the proxy developed by Adler et al (2021) based on changes in central bank reserves adjusted for valuation effects, income flows and changes in other foreign currency balance sheet positions. These data are only available at a monthly frequency.

For the external financing shocks, we use observations of US monetary policy surprises based on high-frequency asset price changes in narrow windows around Federal Open Market Committee (FOMC) announcements from Bauer and Swanson (2023). The surprises have been purged from variation in standard US macroeconomic and financial indicators. They are therefore unlikely to be influenced by common business cycles that would naturally lead to co-movement between US and EME policy rates.

To assess currency mismatch exposure, we rely on variables relating to external portfolio debt liability. Over time, EMEs have increased portfolio debt liabilities relative to GDP and as a share of overall debt financing (Graphs 1.A and 1.B). We use net (rather than gross) liabilities, as foreign assets provide a natural cushion.[9] We distinguish between US dollar-denominated portfolio debt (UPD), capturing the mismatch faced by EME borrowers and therefore the channel related to original sin, and local currency portfolio debt (LPD) held by foreign investors, capturing original sin redux, respectively. While foreign-currency debt still dominates, the local currency share has risen in several countries (Graph 1.C).[10]

EMEs have increased reliance on external funding, led by bond market financing — Graph 1

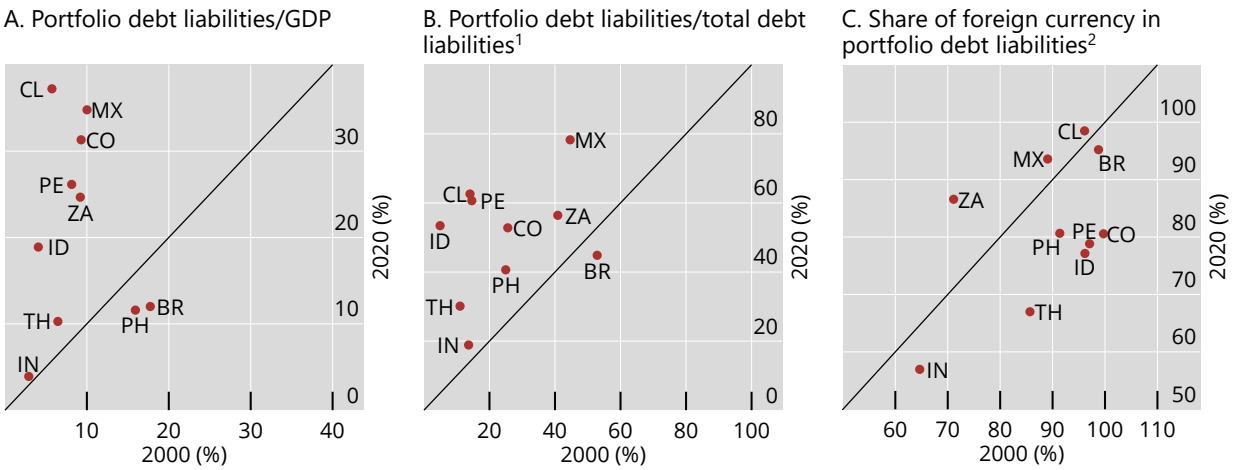

[1] Sum of portfolio debt liabilities and other investment liabilities (such as loans, deposits and trade credits). [2] Assumes the currency composition of portfolio debt liabilities is the same as that of total debt liabilities.

Sources: Allen et al (2023); Lane and Milesi-Ferretti (2018), updated in January 2026.

[9] The cushion, however, can be limited if the entities with the liabilities cannot freely transfer currency risk with the entities holding the assets.

[10] We set aside the credit supply channel given gaps in cross-border bank loan data for certain countries in the initial years of our sample.

We proxy FX market depth by the one-year rolling median bid-ask spread on the FX spot rate (FXS). A higher spread implies less liquid FX markets. Although cruder than order book-based metrics, spread data are easier to obtain for many EME currencies.[11]

Given the centrality of the exchange rate for the overall policy transmission, we also check how the different channels affect the exchange rate itself. For this, we use daily bilateral exchange rates between the United States and the EMEs in our sample. Finally, in our statical analysis we control for exports, imports, GDP growth and inflation.

Our sample comprises 10 EMEs – Brazil, Chile, Colombia, India, Indonesia, Mexico, Peru, the Philippines, Thailand and South Africa – and spans January 2000 to January 2020. These EMEs are classified with free-floating or floating exchange rate regimes.[12],[13] While exchange rates are largely market-determined in these countries, FX interventions can be undertaken in particular situations. The sample ends in January 2020 due to limited availability of currency composition data for external portfolio debt thereafter.

Statistical framework

We use local linear projections to assess the response of EME policy rates to advanced economy monetary policy surprises. The method provides a straightforward statistical way of tracing how a shock propagates to other variables over time. For a given horizon, h, we run a regression of the future change in the outcome – EME policy rates in our case – on the shock:

$$z_{i,t+h} - z_{i,t-1} = \beta_h \cdot mp_t^{AE} + \gamma \cdot x_t + \vartheta_{i,t,h},$$

where $z_{i,t}$ denotes the policy rate of EME i at day t, mp_t^{AE} is the US monetary policy surprise at day t, x_t collects our control variables, which we discuss below, and $\vartheta_{i,t,h}$ is the error term. For each horizon h, the estimate β_h traces how the change in the policy rate responds to a normal US monetary policy surprise. We set $h = 28$ days to allow time for EME policymakers to set the rates following the US announcement.[14]

We build on the baseline specification to test whether countries that have a large currency mismatch or shallow FX markets exhibit higher pass-through of US monetary policy surprises. Specifically, we construct a dummy variable $I_{i,t}^k$ that

[11] As a robustness check, we proxy the ease of managing currency mismatches for each EME currency with FX derivatives turnover and obtain similar results.

[12] Floating regimes are ones where the exchange rate is de facto largely market-determined (Habermeier et al (2009)). If there has been no intervention over the past six months, except for limited intervention to address disorderly market conditions, the regimes are further classified as free-floating.

[13] We start with a broad list of 23 EMEs with reliable daily data on policy rate and exchange rate (Algeria, Argentina, Brazil, Chile, China, Colombia, Hungary, Indonesia, India, Kuwait, Morocco, Mexico, Malaysia, Peru, the Philippines, Poland, Romania, Saudi Arabia, South Africa, Thailand, Turkey, the United Arab Emirates and Vietnam). We then narrow the sample by excluding countries that are classified as floaters or free floaters less than 80% of the time. We also exclude central and eastern European EMEs, which are more closely tied to the euro area.

[14] Mendoza-Fernández and Pelin (2025) document that 80% of non-US central banks already hold their first scheduled policy meeting within 14 days of an FOMC announcement.

indicates if EME i has above-median exposure to measure $k = \text{UPD, LPD, FXS}$ at time t, and interact it with the shock:

$$z_{i,t+h} - z_{i,t-1} = \beta_h^{k,L} \cdot (1 - I_{i,t}^k) \cdot mp_t^{AE} + \beta_h^{k,H} \cdot I_{i,t}^k \cdot mp_t^{AE} + \gamma \cdot x_t + \vartheta_{i,t,h}.$$

where $\beta_h^{k,L}$ is the reaction of the less exposed economies, and $\beta_h^{k,H}$ is the response of the highly exposed economies. We expect $\beta_h^{k,H}$ to be larger than $\beta_h^{k,L}$: when vulnerabilities are higher, central banks are more likely to raise policy rates in step with the US.

We also add controls, x_t. These consist of two additional interaction terms similar to the ones above but for countries with above-median export- or import-to-GDP ratios at time t. This ensures that our estimates are not confounded with trade channel effects. We also include lags of inflation and GDP growth to control for the macroeconomic environment in each EME.[15]

We further augment the specification to test if exposure to both external liabilities and FX spreads simultaneously determine central bank responses, as theory would suggest. In this case, we define indicators, $I_{i,t}^{j,c}$, for the four cases c: low debt, low FX spreads ($L^D L^S$); low debt, high FX spreads ($L^D H^S$); high debt, low FX spreads ($H^D L^S$); and high debt, high FX spreads ($H^D H^S$), and run the regressions:

$$z_{i,t+h} - z_{i,t-1} = \sum_c \beta_{j,c} \cdot I_{i,t}^{j,c} mp_t^{AE} + \gamma \cdot x_t + \vartheta_{i,t,h}.$$

where $j = \text{UPD, LPD}$. In this case, $\beta_{j,c}$ gives the responses in each case.

Policy rate responses

On average, EMEs tend to adjust their policy rates in the same direction as US monetary policy surprises (Graph 2.A, red bar). The pass-through is about 60% (Graph 2.A), similar to that reported by Mendoza-Fernández and Pelin (2025). However, this average response is not statistically significant.

Economies with larger external US dollar-denominated bond portfolios show higher pass-through, roughly moving their policy rates one for one with the US policy rate shock (Graph 2.A). By contrast, countries with lower dependence on US dollar funding do not appear to adjust their policy rates at all to US rate shocks. These patterns are consistent with exchange rate fluctuations playing a greater role in shaping financial conditions in more dollar-dependent economies due to the channel related to original sin, prompting central banks to counter local currency depreciation more actively.

A country's FX market depth also shapes the central bank's response. Countries with less liquid FX markets tend to track the US policy rate more closely than their peers (Graph 2.B). In fact, when we interact external US dollar bond exposure with FX market liquidity, we find that it is countries with *both* high US dollar debt and illiquid FX markets that tend to follow the US policy rate more systematically, with an estimated pass-through approaching twice the US policy rate change (Graph 3.A).[16] By contrast, in countries with higher debt but more liquid FX markets, or with less

[15] Our results are robust to additional controls, including the distance from the capital city to Washington DC, and bilateral trade with the United States (imports and exports) scaled by GDP.

[16] Larger than one-for-one responses are consistent with regimes of imperfect monetary credibility or sudden-stop risk (Montes and Ferreira (2020); Mendoza (2010)), which may necessitate more forceful policy-rate adjustments to anchor inflation or prevent financial turmoil.

EMEs with higher external US dollar debt or less developed FX markets tend to track the US monetary policy rate more closely[1]

In percentage points

Graph 2

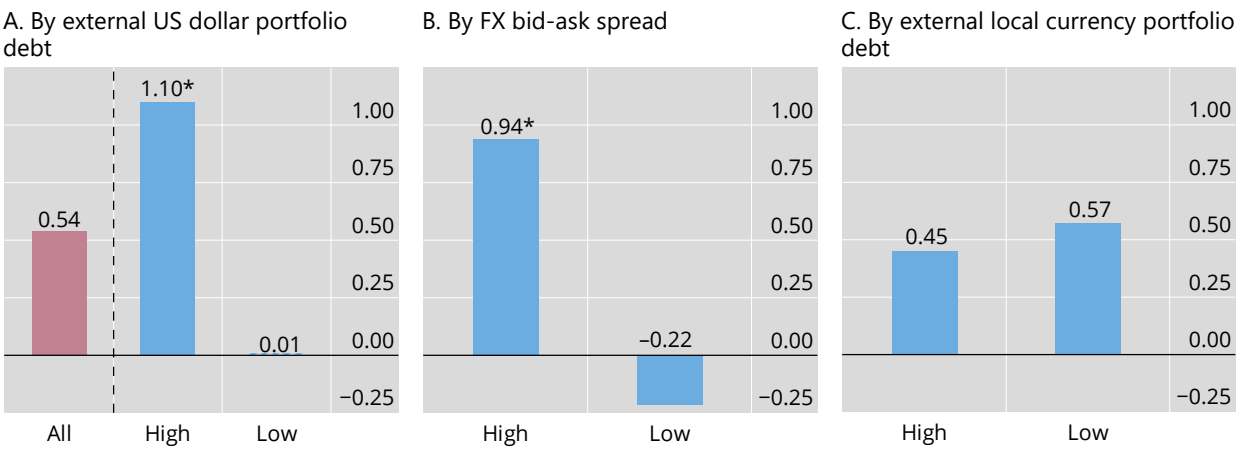

[1] The bars show the change of policy rate after four weeks in response to a US monetary policy surprise corresponding to a 100 basis point increase in the US policy rate. For instance, a 1 percentage point increase in foreign rates increases EME policy rates by approximately 0.5 percentage points on average, while a similar sized easing surprise reduces EME rates by the same amount. ***/**/* denotes statistical significance at the 1/5/10% level. Countries are classified as high/low if the value of the relevant variable is above/below median.

Sources: Allen et al (2023); Bauer and Swanson (2023); Lane and Milesi-Ferretti (2018), updated in January 2026; Bloomberg; authors' calculations.

liquid FX markets but lower debt, central banks do not move rates in tandem with the United States. This pattern suggests complementary roles for external financing and exchange rate market frictions in amplifying the influence of exchange rates on policy responses in EMEs.

Interestingly, larger exposure to externally held local currency-denominated bond portfolios (original sin redux channel) does not go hand in hand with greater pass-through. Specifically, both high and low (LPD) countries have pass-throughs that are close to the panel average and statistically insignificant (Graph 2.C). The difference between them is also insignificant. One possible explanation is that global investors' appetite for a given EME is closely linked to their appetite for EME assets as a whole, which tends to wax and wane with the broad dollar index rather than any bilateral exchange rate. Because an individual EME can at most influence its own bilateral rate, the policy payoff from raising rates to retain global investor interest is probably limited. Accordingly, in these cases EME central banks may place less priority on curbing exchange rate fluctuations.

FX interventions

FX intervention offers an additional tool – alongside interest rate adjustments – for responding to foreign monetary policy surprises. By selling FX reserves to support the domestic currency, central banks can lean against excessive local currency depreciation, mitigating the financial tightening associated with weaker exchange rates. How desirable and effective interventions are, however, varies across circumstances. For example, countries with shallow FX markets may have a greater need to provide liquidity to limit excessive exchange rate volatility in times of adverse

Countries with high external US dollar debt combined with less developed FX markets see stronger policy responses and smaller currency fluctuations[1] Graph 3

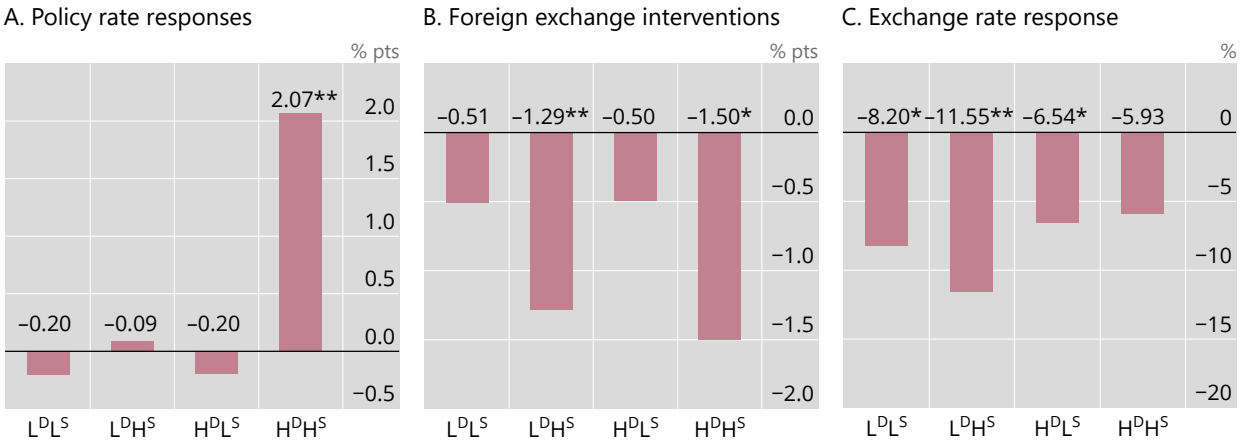

L^D/H^D = low/high external US dollar debt; L^S/H^S = low/high FX market spreads.

[1] The bars show the response of the policy rate, foreign exchange intervention and exchange rate after four weeks to a US monetary policy surprise corresponding to a 100 basis point increase in the US policy rate. ***/**/* denotes statistical significance at the 1/5/10% level. Countries are classified as high/low if the value of the relevant variable is above/below median.

Sources: Adler et al (2023); Allen et al (2023); Bauer and Swanson (2023); Lane and Milesi-Ferretti (2018), updated in January 2026; Bloomberg; authors' calculations.

external shocks.[17] Moreover, in such markets intervention can be relatively effective at improving market liquidity and is likely to deliver a larger price impact for a given intervention size in the short run.

We apply our statistical framework to assess FX intervention responses to US monetary surprises, replacing the change in the policy rate with our measure of FX intervention. As the data are monthly, we focus on intervention in the month after each FOMC announcement.

We find that EMEs, on average, deploy FX intervention in response to US monetary policy surprises (Graph 4.A). They tend to sell FX after policy tightening (to counter local currency depreciation) and buy FX after easing (to counter appreciation).

This systematic reaction is most prominent in countries with thin FX markets. Responses in such countries are more than twice as large relative to their peers – and statistically significant (Graph 4.B). In contrast to the policy rate response, whether a country has high or low external debt, including in US dollars, does not appear to materially affect central banks' decision to intervene (Graphs 4.A and 4.C; Graph 3.B). These results suggest that central banks consider not only the need but also the desirability and expected effectiveness of intervention when deploying it to ensure orderly market functioning and limit excessive exchange rate fluctuations.

[17] Using a recent survey of 21 EME central banks, Cavallino and Patel (2019) document that restricting exchange rate volatility and providing market liquidity are seen as important immediate objectives. The central banks find that FX interventions are mostly effective for up six months, but not beyond.

EMEs with less developed FX markets tend to deploy larger FX interventions following US monetary policy surprises[1]

In percentage points

Graph 4

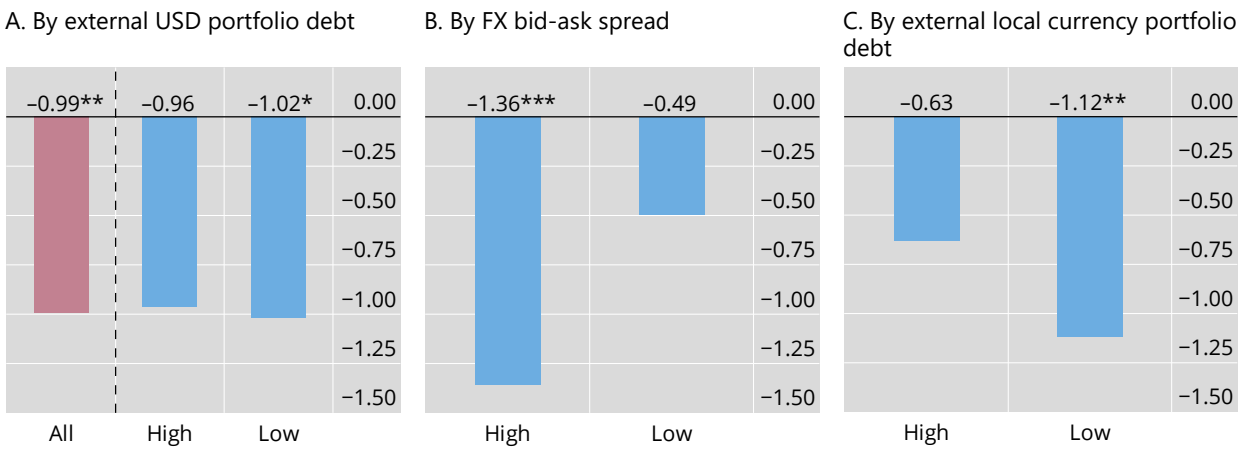

[1] The bars show the FX intervention proxied by the adjusted change in foreign reserves as a share of GDP after one month in response to a US monetary policy surprise corresponding to a 100 basis point increase in the US policy rate. For instance, a 1 percentage point increase in foreign rates decreases EME reserves by approximately 1 percentage point in terms of GDP on average, while a similar sized easing surprise increases EME reserves by the same magnitude. ***/**/* denotes statistical significance at the 1/5/10% level. Countries are classified as high/low if the value of the relevant variable is above/below median.

Sources: Adler et al (2023); Allen et al (2023); Bauer and Swanson (2023); Lane and Milesi-Ferretti (2018), updated in January 2026; authors' calculations.

Exchange rate responses

Exchange rate movements are central to the way EME central banks respond to foreign policy rate surprises through the external financing channel. Our previous results strongly support the notion that EME central banks mitigate the resulting contractionary or expansionary effects of FX moves by using policy rate adjustments and FX intervention. We now study the overall effects on exchange rates themselves, after the deployment of these tools. To do so, we use our empirical framework with the 28-day change in exchange rates on the left-hand side, ie the same horizon at which we assess the policy responses.

EMEs with larger external US dollar debt and thinner FX markets than their peers show equal or slightly smaller exchange rate pass-through after US monetary surprises. Although all EME currencies typically depreciate after a US tightening and appreciate after an easing, we do not observe economically and statistically significant differences between countries (Graphs 5.A and 5.B). Indeed, among countries with both higher external US dollar debt and less liquid FX markets, the average depreciation is the smallest and not statistically significant (Graph 3.C). In contrast, countries with higher external local currency debt tend, on average, to experience more than twice the depreciation (Graph 5.C). These economies are more exposed to the external financing shocks, and their central banks typically neither adjust policy rates nor deploy FX intervention in response to US monetary policy shocks. As a result, they show larger exchange rate adjustment than their peers.

One interpretation of these results is that central banks' response to US monetary policy surprises, either through rate adjustment or FX interventions, helps dampen

EMEs, especially the ones with higher external local currency debt, experience larger currency depreciation after US monetary policy surprises[1]

In per cent

Graph 5

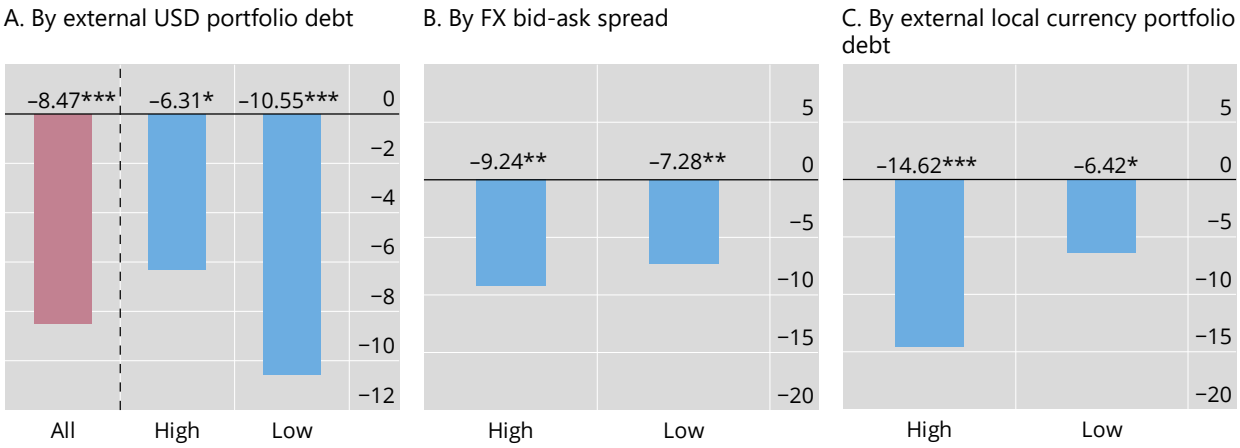

[1] The bars show the log change in exchange rate (quoted as local currency per dollar) after four weeks in response to a US monetary policy surprise corresponding to a 100 basis point increase in the US policy rate. For instance, a 1 percentage point increase in foreign rates approximately leads to a 10% depreciation of EME currencies on average, while a similar sized easing surprise leads to a 10% appreciation of EME currencies. ***/**/* denotes statistical significance at the 1/5/10% level. Countries are classified as high/low if the value of the relevant variable is above/below median.

Sources: Adler et al (2023); Allen et al (2023); Bauer and Swanson (2023); Lane and Milesi-Ferretti (2018), updated in January 2026; Bloomberg; authors' calculations.

exchange rate movements. Absent such responses, these economies would probably face sharper tightening of financial conditions in response to external financing shocks, amplifying local currency depreciation. The similarity in exchange rate outcomes across exposure groups suggests that countervailing policies have mitigated spillovers.

Conclusion

In this article, we investigate how EME central banks respond to US monetary policy shocks. We find that EMEs with large foreign currency liabilities and shallow FX markets tend to adjust policy rates in the same direction as the US shock. When market depth is limited, these economies supplement the rate reactions with FX intervention to provide liquidity and limit excessive and potentially destabilising exchange rate movements. Other economies do not respond strongly on average.

This policy mix appears to dampen excessive currency volatility in vulnerable economies. Exchange rates move more strongly in economies with lower external debt or deeper FX markets in response to US monetary policy shocks. The relative attenuation of exchange rate movements in vulnerable economies is consistent with the policy mix leaning against balance sheet amplification arising from unhedged foreign currency exposures.

Our results highlight that the transmission of monetary policy can change if financial and FX market frictions become too large. Mapping the turning points when this happens can help central banks to identify when to adjust the policy mix, and

where structural measures – such as deepening local FX markets and promoting hedging – yield the largest dividends. A better understanding of these mechanisms can also aid market participants in interpreting policy choices, improving expectations formation and price discovery.

References

Adler, G, K S Chang, R Mano and Y Shao (2021): "Foreign exchange Intervention: a dataset of public data and proxies", *IMF Working Paper*, no 21/47.

Akıncı, Ö and A Queralto (2024): "Exchange rate dynamics and monetary spillovers with imperfect financial markets", *The Review of Financial Studies*, vol 37, no 2.

Allen, M, D Gautam, and L Juvenal (2023): "Currencies of external balance sheets", *IMF Working Papers*, no 237.

Bank for International Settlements (BIS) (2019): *Annual Economic Report 2019*, Chapter II.

——— (2023): *Inflation, external financial conditions and macro-financial stability frameworks in Asia-Pacific*, October.

Basu, S, E Boz, G Gopinath, F Roch and D Unsal (2025): "Integrated monetary and financial policies for small open economies", *Econometrica*, vol 93, no (6).

Bauer, M and E Swanson (2023): "A reassessment of monetary policy surprises and high-frequency identification", *NBER Macroeconomic Annual* 23.

Brandão-Marques, L, G Gelos, M Narita and E Nier (2021): "Leaning against the wind: An empirical cost-benefit analysis", *CEPR Discussion Paper Series*, no 15693.

Bruno, V and H S Shin (2015): "Cross-border banking and global liquidity", *Review of Economic Studies*, vol 82, no 2.

Calvo, G and C Reinhart (2002): "Fear of floating", *The Quarterly Journal of Economics*, vol 117, no 2.

Carstens, A and H S Shin (2019): "Emerging markets aren't out of the woods yet", *Foreign Affairs*, 15 March.

Cavallino, P and N Patel (2009): "FX intervention: goals, strategies and tactics", *BIS Papers*, no 104.

Céspedes, L F, R Chang and A Velasco (2004): "Balance sheets and exchange rate policy", *American Economic Review*, vol 94, no 4.

Checo, A, F Grigoli and D Sandri (2024): "Monetary policy transmission in emerging markets: proverbial concerns, novel evidence", *BIS Working Papers*, no 1170.

Claessens, S, L Stracca and F Warnock (2016): "International dimensions of conventional and unconventional monetary policy", *Journal of International Money and Finance*, vol 67.

Cook, D (2004): "Monetary policy in emerging markets: Can liability dollarization explain contractionary devaluations?", *Journal of Monetary Economics*, vol 51.

De Leo, P, G Gopinath and S Kalemli-Özcan (2022): "Monetary policy and the short-rate disconnect in emerging economies", *NBER Working Papers*, no 30458.

Eichengreen, B and R Hausmann (1999): "Exchange rates and financial fragility", *NBER Working Papers*, no 7418.

Gelos, G and R Sahay, eds (2023): *Shocks and capital flows: policy responses in a volatile world*, International Monetary Fund.

Gopinath, G, E Boz, C Casas, F J Díez, P-O Gourinchas, and M Plagborg-Møller (2020): "Dominant currency paradigm", *American Economic Review*, vol 110, no 3.

Gourinchas, P-O (2018): "Monetary policy transmission in emerging markets: an application to Chile", Central Bank of Chile *Central Banking, Analysis, and Economic Policies Book Series*, vol 25.

Habermeier, K, A Kokenyne, R Veyrune and H Anderson (2009): "Revised system for the classification of exchange rate arrangements", *IMF Working Papers*, no 211.

International Monetary Fund (IMF) (2020): "Toward an integrated policy framework", *Policy Papers*, no 046.

Juselius, M, P Wooldridge and D Xia (2025): "Financial channel implications of a weaker dollar for emerging market economies", *BIS Bulletins*, no 114.

Kearns, J, A Schrimpf, and D Xia (2023): "Explaining monetary spillovers: The matrix reloaded", *Journal of Money, Credit and Banking*, vol 55.

Lane, P and G M Milesi-Ferretti (2018): "The external wealth of nations revisited: international financial integration in the aftermath of the global financial crisis", *IMF Economic Review*, vol 66.

Mano, R C and S Sgherri (2024): "One shock, many policy responses", *Open Economies Review*, vol 35, no 2.

Mendoza, E G (2010): "Sudden Stops, Financial Crises, and Leverage", *American Economic Review,* vol 100, no 5.

Mendoza-Fernández, A and S Pelin (2025): "Do central banks follow the Fed?", working paper.

Montes, G C and C F Ferreira (2020): "Does monetary policy credibility mitigate the fear of floating?", *Economic Modelling*, vol 84.

Shin, H S, P Wooldridge and D Xia (2025): "US dollar's slide in April 2025: the role of FX hedging", *BIS Bulletin*, no 105.

Iñaki Aldasoro
inaki.aldasoro@bis.org

Bryan Hardy
bryan.hardy@bis.org

Goetz von Peter
goetz.vonpeter@bis.org

Philip Wooldridge
philip.wooldridge@bis.org

International finance through the lens of BIS statistics: offshore activity[1]

Offshore activity has become an integral part of international finance, with companies increasingly issuing bonds via financial affiliates abroad while banks and other financial firms channel credit through financial centres. This article leverages BIS statistics to look through residence to nationality to show who ultimately borrows and lends offshore. By revealing who borrows and which banking systems hold the exposures, BIS statistics enhance the monitoring of vulnerabilities and international spillovers.

JEL classification: F23, F36, G15.

The conventional framework for macroeconomic analysis aligns economic and financial activity with residence and country borders. Yet multinational borrowers and creditors often transact through affiliates in financial centres abroad rather than from offices in their home market. Such offshore activity represents a sizeable share of financial activity. While enhancing financial integration, offshore activity also obscures international linkages in conventional residence-based statistics.

This article is a primer on how the BIS international banking and financial statistics can be used to analyse offshore financial activity. Offshore activity is commonly associated with indirect flows between borrowers and savers through complex corporate structures straddling borders. Several BIS data sets, especially the international debt securities (IDS) and international banking statistics (IBS), help to shed light on offshore activity by looking through the sector and residence of an entity to those of its parent, defined by its corporate headquarters. In other words, they provide a *nationality* view that complements the conventional *residence* perspective by reallocating the balance sheets of foreign affiliates to their parents.

The nationality view has a long history in BIS statistics and analysis. It had a foundational influence on the design and evolution of BIS international banking and financial statistics (CGFS (2000); Borio (2013)). Reallocation by nationality later formed the basis for broader research efforts to examine who drives economic activity, how

[1] The views expressed in this publication are those of the authors and not necessarily those of the BIS or its member central banks. This is the eighth article in a series showcasing the BIS international banking and financial statistics. We thank Stefan Avdjiev, Torsten Ehlers, Jon Frost, Carlos Mallo, Patrick McGuire, Benoît Mojon, Swapan-Kumar Pradhan, Daniel Rees, Andreas Schrimpf, Hyun Song Shin, Frank Smets and Sonya Zhu for helpful comments and Jeemin Son for excellent research assistance.

Key takeaways

- *Offshore activity obscures financial links in residence-based statistics that only become clear from a nationality perspective, like that available in BIS statistics on banking and bond markets.*
- *Almost a third (nearly $11 trillion) of international debt securities (IDS) are issued by affiliates located outside their home country. Non-financial companies issue nearly half of IDS via non-bank financial affiliates, often in financial centres, with proceeds sometimes repatriated via intercompany loans.*
- *Banks' offshore affiliates, mainly in financial centres, account for over 40% of global cross-border bank lending. Corresponding statistics for portfolio investment and FDI on a nationality basis are lacking.*

corporate groups are interconnected and how financial vulnerabilities and spillovers might arise (Avdjiev et al (2014); Coppola et al (2021); Beck et al (2024)).

This primer extends the residence vs nationality analysis in McGuire et al (2024) by tracing the sectors, instruments and economies through which offshore activity is reallocated in banking and bond markets. It highlights how offshore borrowing can alter assessment of financial vulnerabilities and international spillovers when viewed from a nationality perspective, as they are obscured under the conventional residence view. It shows the extent to which non-financial corporations (NFCs) issue bonds through their foreign financial affiliates – often incorporated in financial centres – and repatriate the proceeds as intercompany loans. The residence view complicates assessments of a debtor's vulnerability to shocks that originate from foreign creditors by masking the nationality of the creditor.

The first section discusses the concept of offshore activity. The second dissects offshore debt securities issuance. The subsequent section turns to offshore lending and investment. The conclusion sketches challenges that offshore activity and related data gaps pose for financial stability analysis.

The concept of offshore activity

Offshore activity refers to financial activity abroad that is unrelated to developments in the economy where the activity takes place.[2] Offshore activity takes many forms, from bond issuance in foreign markets to investment funds that pool non-residents' savings and currency trading beyond the reach of regulations that apply to transactions onshore. What these forms have in common is that the activity is ultimately driven by agents that are not residents of the jurisdiction in which it takes place.[3] In that sense, offshore activity is inherently cross-border.

Offshore activity concentrates in financial centres. These range from large global centres that provide a full range of financial services, like London and New York, to those specialised in certain services and located in small economies, like the Cayman Islands and Luxembourg (Box A). Economies of scale and scope benefit global

[2] A narrower notion of offshore activity emerged from the rise of eurocurrency markets in the 1950s and 1960s, where the defining characteristic is that activity is denominated in a currency that is foreign to both counterparties (McCauley et al (2021)).

[3] That is, offshore activity excludes business driven by the domestic supply of or demand for funds.

centres. At the same time, physical distance, regulation and taxation work against the tendency of financial activity to concentrate (Pogliani et al (2022)). This creates opportunities for multiple financial centres to co-exist even within time zones.

Graph 1 illustrates how offshore activity via financial centres complicates the interpretation of capital flows and external assets and liabilities. Consider a multinational firm that raises funding through an affiliate located in a financial centre. Such affiliates are often set up as special purpose vehicles (SPVs) or other types of non-bank financial institutions (NBFIs) (Graph 1, red boxes). For example, Vale, a Brazilian mining company, issues bonds through an SPV domiciled in the Cayman Islands. The proceeds might finance international operations, or be channelled to Vale's home office via an intercompany loan (red arrows). In the latter case, the issuance would be recorded in conventional balance of payments statistics as debt incurred by residents of the financial centre, while the intercompany loan would be recorded as foreign direct investment (FDI) from the financial centre to the firm's home country (ie from the Cayman Islands to Brazil).

Banks and NBFIs also use financial centres for borrowing and lending. Instead of lending directly to an NFC abroad from its head office in country A (Graph 1, black arrow at the bottom), a bank could lend indirectly via a foreign affiliate, funding the loan from its head office via internal capital markets (purple arrows). Such an indirect loan would be recorded twice: as external debt of the financial centre where the affiliate resides and as external debt of country B where the NFC is based. NBFIs similarly channel funds through and within financial centres (green boxes). For example, many hedge funds domiciled in the Cayman Islands hold US Treasuries to profit from the cash-futures basis trade (Barth et al (2025)), though most are owned or financed by entities in the United States (Bertaut et al (2021)).

The residence perspective on international finance is based on a "triple coincidence": the GDP area, currency area and location of decision-making units (and their balance sheets) are assumed to coincide with country borders (Avdjiev et al (2016)). When these align, financial transactions between residents and non-residents are driven by domestic savings and investment and subject to domestic laws and regulations. However, as Graph 1 illustrates, offshore activity multiplies linkages and distorts observed financial flows. It obscures the underlying drivers of

A stylised illustration of offshore activity — Graph 1

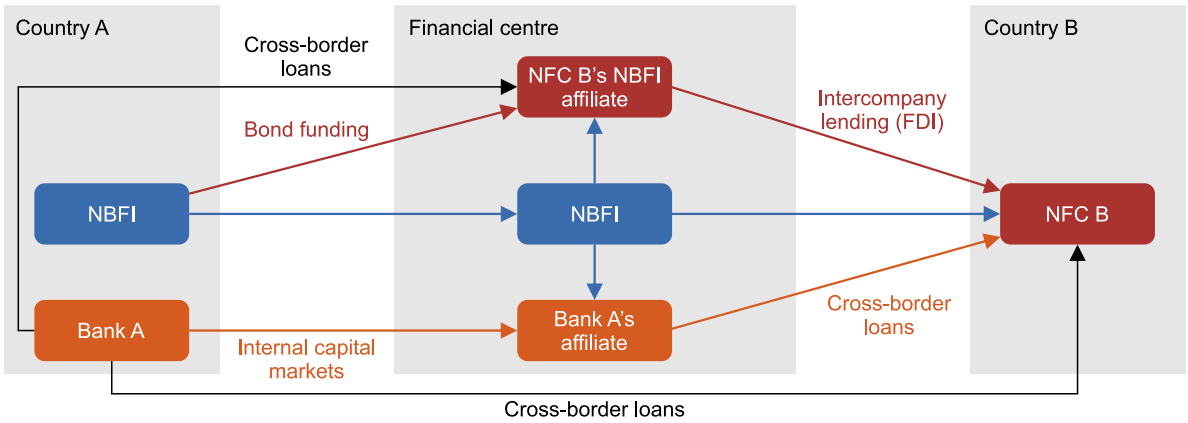

NBFI = non-bank financial institution; NFC = non-financial corporation.

Source: Authors' elaboration.

borrowing and lending, and changes the economic interpretation of financial flows – such as FDI that behaves more like debt (Blanchard and Acalin (2016)).

Offshore activity calls for an approach that consolidates activity along lines of ownership and control. The resulting nationality perspective recognises that decision-making units straddle country borders when firms are multinational, breaking the triple coincidence. On the borrower side, linking assets and liabilities to the ultimate parent identifies the country and sector that control them. For example, bonds issued by Vale's subsidiary represent liabilities controlled by its Brazilian parent and have little relation to economic activity in the Cayman Islands. On the lender side, reallocating bank positions by nationality uncovers which institutions ultimately bear credit exposures and funding shortages (McGuire and von Peter (2009)).

The nationality view complements the residence view by providing insights into who drives decisions affecting international assets and liabilities. However, to assess who is ultimately responsible for repaying debts, information about nationality is insufficient. The extent to which a borrower's liabilities are backed by its parent depends on corporate structures and guarantees (Box B).

Offshore activity and the nationality view have a long history in BIS statistics. The initial collection of the locational banking statistics, in the 1960s, was motivated by the growth of "eurocurrency" markets, where banks borrowed and lent foreign currency from outside their home market, mostly in US dollars (McCauley et al (2021)). The expansion of offshore banking later led to the collection of the consolidated banking statistics, which put the nationality of banks at the core. The growth of the "eurobond" market in the 1980s led to the introduction of the IDS. Being compiled from data on individual securities, they have been aggregated from inception by the residence of the immediate issuer as well as the nationality of the issuer's parent (Gruić and Wooldridge (2012)).

Unmasking the offshore activities of bond issuers

One of the markets shaped by offshore activity is the international bond market. Of the $33 trillion in international bonds outstanding at end-Q3 2025, almost $11 trillion was issued by affiliates located offshore, incorporated in countries other than where their parents were headquartered (Graph 2.A).[4] Over the past three decades, issuance by such offshore affiliates has grown faster than aggregate issuance of IDS, with its share rising from around 20% in the mid-1990s to 31% in 2025.

Issuance by offshore affiliates is highly concentrated in financial centres. Global financial centres, notably London and New York, unsurprisingly play a role given their size and markets' liquidity (Graph 2.B). However, issuance is much larger from affiliates in so-called cross-border financial centres (XFCs) that cater to non-residents, like the Cayman Islands, Ireland and Luxembourg (Box A). Issuance by foreign affiliates residing in these centres is considerably larger than that from all other

[4] In order to classify an issue in the IDS statistics, the BIS assesses: (i) the residence of the immediate issuer; (ii) the location of the issue's registration; (iii) the governing law; and (iv) the listing location. When all four characteristics refer to the same country, the issue is classified as a domestic debt security. When at least one points to a different country, it is instead classified as international. See Box A in Aldasoro et al (2021) for further details.

A rising share of international debt securities is issued by foreign affiliates

Graph 2

A. IDS issued by affiliates located onshore in the home country and offshore in other countries

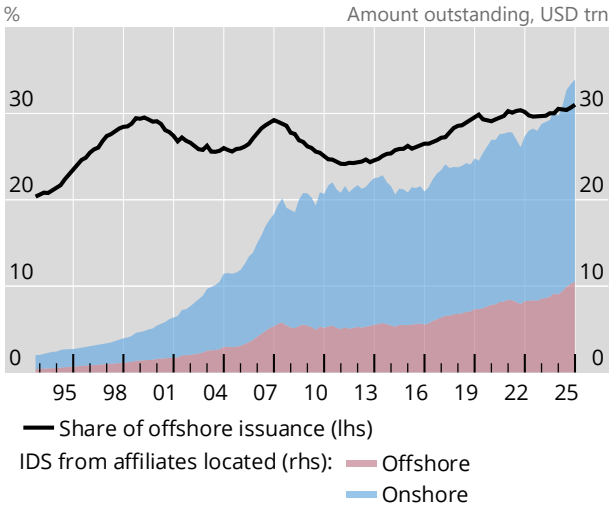

B. IDS issued by foreign affiliates, by residence and sector of the immediate issuer[1]

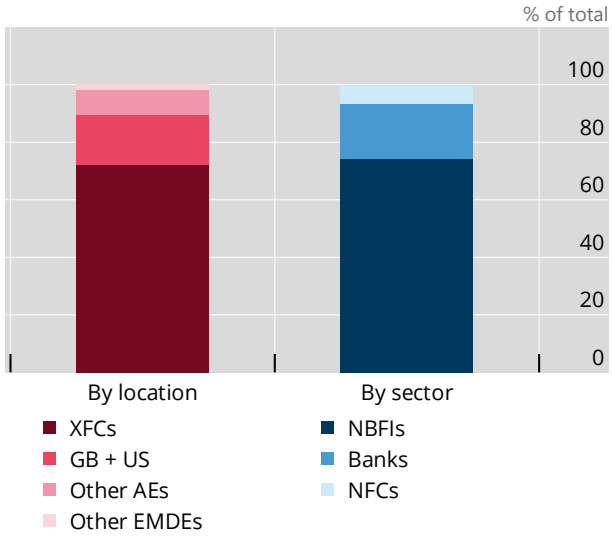

IDS = international debt securities (excludes domestic debt securities); XFCs = cross-border financial centres (see Box A); AEs = advanced economies; EMDEs = emerging market and developing economies; NBFIs = non-bank financial institutions; NFCs = non-financial corporations.

[1] Based on IDS outstanding at end-Q3 2025 issued by affiliates located outside the home country (the red area in Graph 2.A), excluding IDS issued by international organisations ($3.0 trillion).

Sources: Dealogic; Euroclear; LSEG; Xtrakter Ltd; authors' calculations.

locations combined. Most of these foreign affiliates are NBFIs, although for many the sector of the immediate issuer is different from that of the parent.[5]

A comparison of IDS aggregated by nationality and sector of the parent versus by residence and sector of the immediate issuer illustrates how offshore activity creates a wedge between residence and nationality statistics. Graph 3 plots, for a given sector and country group, the ratio between IDS amounts outstanding by nationality and by residence. For example, for emerging market economy (EME) banks (Graph 3.A), a ratio of two indicates that IDS issued by banks with *parents* in EMEs are twice as large as the amounts owed by banks *located* in EMEs. Conventional residence-based statistics thus underestimate the consolidated debt of EME banks.

Nationality- and residence-based debt measures vary considerably across countries and sectors. For advanced economy (AE) banks, IDS issuance is similar on a nationality and on a residence basis (Graph 3.A, red line close to one). For EMEs, in contrast, the ratio shot up over the past 20 years, with Chinese banks playing a prominent role (solid and dashed blue lines). At the other extreme, for financial centres the ratio is below one, in line with the role they play in offshore activity (black line). The picture is broadly similar, though slightly less pronounced, for NBFIs (Graph 3.B). Notably, the ratio for financial centres is closer to zero than to one, since much of the issuance of NBFIs is by affiliates of foreign entities operating there.

Issuance of debt securities by offshore affiliates is most prevalent among NFCs. In general, NFCs borrow through offshore affiliates to tap deeper, diverse investor bases, raise larger amounts in longer maturities, align funding currencies with

[5] Governments, the largest players in the global bond market, do not have offshore affiliates and mostly issue local currency bonds through their domestic debt offices.

IDS are larger on a nationality than a residence basis in many economies

Ratio of IDS outstanding by nationality to those by residence[1] Graph 3

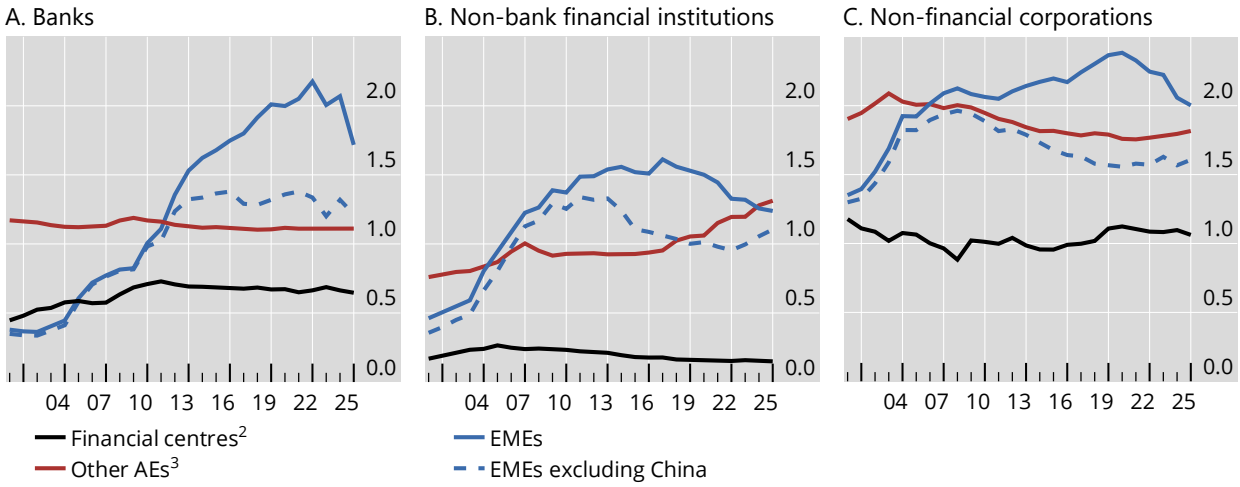

[1] International debt securities (IDS) issued by entities whose ultimate parent is a bank (panel A), NBFI (panel B) or NFC (panel C) headquartered in the specified country group divided by IDS issued by entities whose immediate issuer is a bank (panel A), NBFI (panel B) or NFC (panel C) residing in that country group. [2] See Box A. [3] Advanced economies excluding cross-border financial centres.

Sources: Dealogic; Euroclear; LSEG; Xtraker Ltd; authors' calculations.

revenues and, in some cases, benefit from established international legal and market infrastructures. For AEs, the nationality-to-residence ratio in bond issuance is structurally high (Graph 3.C, red line). For EMEs, it expanded sharply after 2000 (blue lines), driven by Brazil, Russia and especially China. Excluding China, the ratio for EME NFCs is now similar to that of AE NFCs. Chinese NFCs continue to issue more through their offshore affiliates than other EMEs, even as volumes declined recently.

Restating bond issuance on a nationality basis also reveals the *sector* of the ultimate corporate parent – a reallocation that is particularly sizeable for NFCs. On an ultimate parent basis, NFCs are much more important borrowers than their issuance on an immediate borrower basis suggests. As part of their strategy for optimising funding costs and taxes, NFCs often raise funding through financial affiliates, typically NBFI subsidiaries (eg SPVs) domiciled abroad. This is more common among EME corporates due to capital account restrictions and less developed domestic markets. Of $6.4 trillion in IDS owed by AE-headquartered NFCs at end-Q3 2025, $2.6 trillion (41%) was issued by NBFI affiliates (Graph 4.A). For EME-headquartered NFCs, $0.7 trillion (48%) out of $1.5 trillion was issued by NBFI affiliates (Graph 4.B).

Offshore affiliates of NFCs are located in a handful of financial centres. Euro area NFCs issue relatively more from affiliates in European financial centres such as Luxembourg and the Netherlands (Graph 5.A). US NFCs rely more on affiliates in the Cayman Islands, whereas other AE NFCs (eg Swiss, UK and Japanese) issue through affiliates in a wider set of jurisdictions.

NBFI issuers are often affiliates of parents from other sectors[1]

IDS outstanding as of end-Q3 2025, in billions of US dollars — Graph 4

A. Issuers from advanced economies

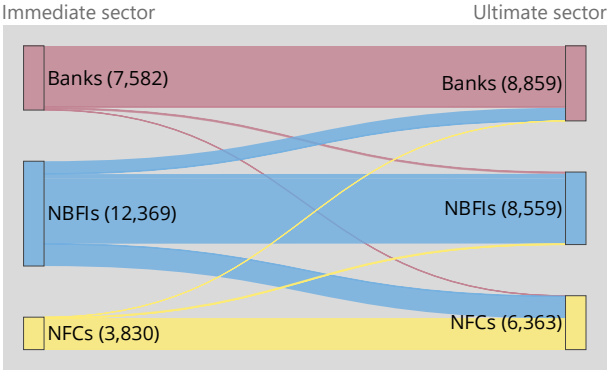

B. Issuers from emerging market economies

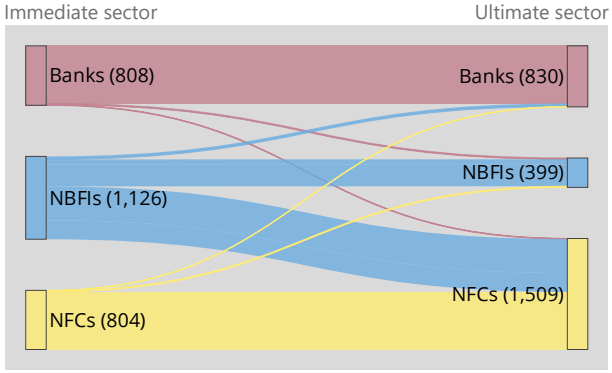

NBFIs =non-bank financial institutions; NFCs = non-financial corporations.

[1] Immediate sector (lhs) corresponds to the sector of the affiliate that issued the bonds. Ultimate sector (rhs) corresponds to the sector of the parent, headquartered in advanced economies (panel A) or emerging market economies (panel B). Numbers in brackets refer to international debt securities (IDS) outstanding by nationality, excluding issuance by governments and international organisations.

Sources: Aldasoro et al (2021); Dealogic; Euroclear; LSEG; Xtrakter Ltd; authors' calculations.

For EME NFCs, the patterns are somewhat starker. Chinese corporates stand out, issuing similar amounts through offshore affiliates as all other EMEs combined (Graph 5.B). Their offshore bond issuance is largely conducted through Caribbean centres, notably the British Virgin Islands and the Cayman Islands. Other EMEs such as Brazil and South Africa, by contrast, tend to borrow through offshore affiliates in European hubs, particularly the Netherlands and Luxembourg.

Where are foreign affiliates of corporate bond issuers located?[1]

IDS issued by foreign affiliates of non-financial corporations, in billions of US dollars — Graph 5

A. Issuers from selected advanced economies

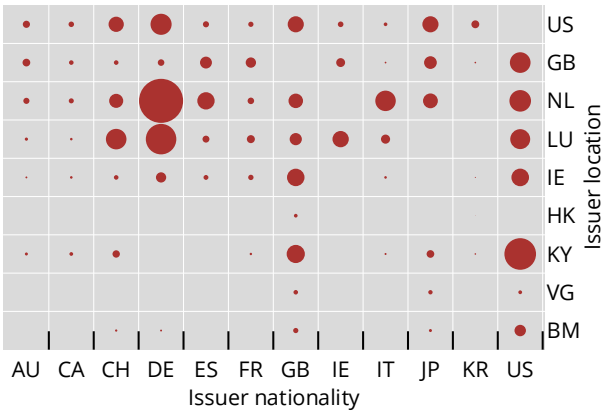

B. Issuers from selected emerging market economies

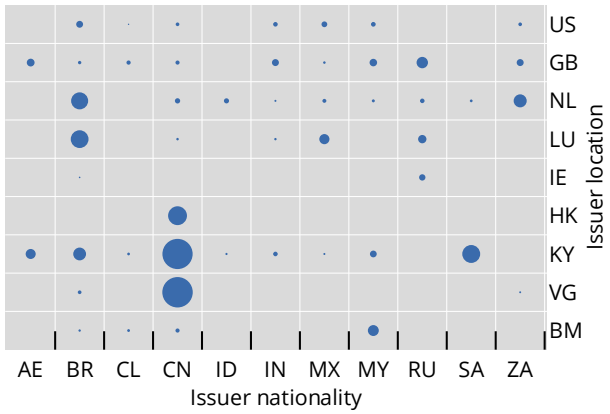

[1] Bubble size is proportional to amounts outstanding at end-Q3 2025. Excludes debt securities issued in domestic markets, eg bonds issued in the United States by US affiliates of foreign companies. IDS = international debt securities.

Sources: Dealogic; Euroclear; LSEG; Xtrakter Ltd; authors' calculations.

Foreign affiliates can repatriate IDS proceeds as intercompany lending

Net flows, in billions of US dollars

Graph 6

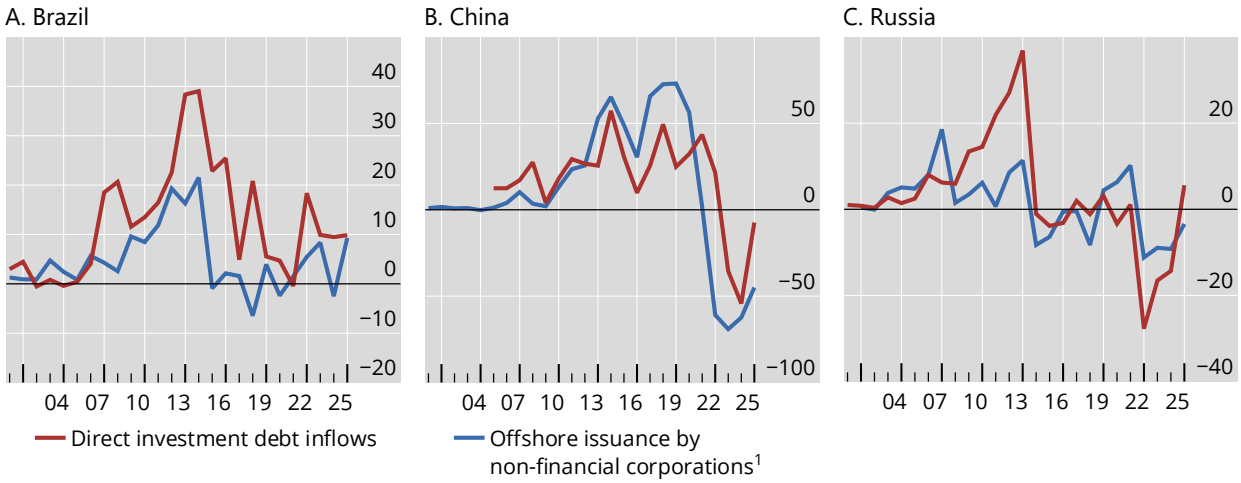

[1] Net issuance of international debt securities (IDS) by foreign affiliates of non-financial corporations headquartered in the specified country.

Sources: IMF; Dealogic; Euroclear; LSEG; Xtrakter Ltd; authors' calculations.

Bond issuance through offshore affiliates obscures the true nature of countries' external accounts in ways that depend on how the funds are deployed across the corporate structure. Evidence of funds from offshore issuance being transferred back to headquarters is most compelling for the largest EMEs (Avdjiev et al (2014)). Comparing offshore issuance with direct investment debt inflows to the home country reveals a strong correlation for Brazil, China and Russia (Graph 6), indicating that intercompany debt inflows tend to rise in periods when more debt is issued by offshore affiliates and fall when new issuance slows. As described in the previous section, this creates two external account entries for one issuance: a portfolio debt liability inflow into the financial centre, and an FDI liability inflow into the home country (with a corresponding FDI asset outflow from the financial centre).

Reallocating offshore lending and investing

Understanding a debtor's vulnerability to shocks originating abroad requires knowing who its creditors are. For example, a shock to the consolidated balance sheet of the lending bank might lead it to reduce lending to borrowers, including unaffected ones, in order to restore its capital adequacy ratio, meet margin calls or adhere to value-at-risk or similar models (Kaminsky and Reinhart (1999)). Data by residence mask the identity of some creditors when lending and investing takes place offshore, routed via intermediaries in financial centres. Residence data thus obscure a debtor's vulnerability to international spillovers via creditors. The nationality perspective in the IBS sheds light on this issue, at least for intermediation via banks.

Bank credit intermediated through financial centres represents a large share of overall bank cross-border credit. Credit extended from banks' home offices accounts for less than 60% of their cross-border claims; the remainder is extended via their foreign affiliates, located mainly in financial centres. The United Kingdom is the largest centre for such offshore banking activity, followed by Hong Kong SAR

(Graph 7.A, blue bars). Bank business models, as well as differences in regulation and tax treatment, are important factors in banks' choice to operate in financial centres.[6]

Looking through credit intermediated through financial centres materially changes the source of bank credit for borrower countries. Borrowers in many countries rely on cross-border borrowing from banks located in financial centres (Graph 7.B, black dots), most of which are headquartered elsewhere. For some borrower countries, banks headquartered in the very same country are among the largest cross-border creditors. Over 40% of cross-border bank credit to borrowers in mainland China is extended by the offices of Chinese banks in financial centres outside the mainland, mainly in Hong Kong SAR (brown bars). Similarly, a quarter of cross-border bank credit to Brazil is from Brazilian bank offices in financial centres outside Brazil, mainly in the Cayman Islands and the Bahamas.

Whereas offshore activity (and the attendant distortions in residence-based statistics) was once synonymous with banking activity, over the past few decades NBFIs have expanded their presence offshore. For example, many asset management companies have incorporated funds in financial centres to invest in bonds and equities globally. Unlike for banks, information about who owns and controls the portfolios of NBFIs is not readily available. BIS statistics provide some information about banks' foreign bond holdings, but their share of total foreign bond holdings is

Banks often lend through their offshore affiliates in financial centres

Outstanding cross-border bank claims, at end-Q3 2025 — Graph 7

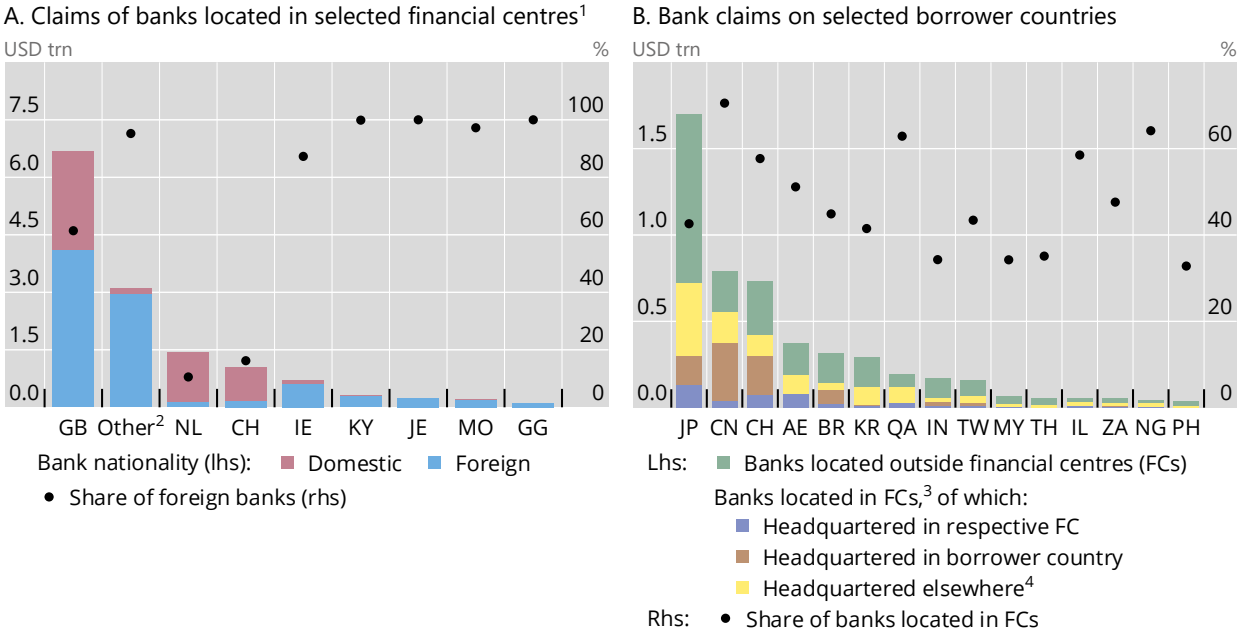

[1] Domestic banks are those headquartered in the financial centre. [2] Includes BM, HK, IM and LU. [3] Financial centres shown in panel A. [4] Includes banks located in one financial centre but headquartered in a different one, as well as partial data for banks headquartered in borrower countries that are not BIS reporting countries.

Sources: BIS locational banking statistics by nationality; authors' calculations.

[6] Banks that adopt a centralised model typically channel a sizeable portion of their international business via financial centres, often through branches. Conversely, banks that follow a decentralised model operate primarily through locally incorporated and capitalised subsidiaries and thus have less of an offshore presence (Hardy et al (2024)).

small. BIS banking statistics can shed light on the scale and geographical distribution of NBFIs' activities only through their links with banks (Box C). However, these links are recorded on a residence basis and do not identify NBFIs' nationality.

Conclusions

Offshore activity complicates financial stability analysis by obscuring links between related entities because multinational firms and banks borrow and lend through their affiliates worldwide. BIS statistics help to reveal those links by looking through the geography of bond issuance (in IDS) and bank credit (in IBS) to the nationality of the entities involved. Reallocating offshore activity by nationality thus highlights financial vulnerabilities that might not be evident from residence data alone.

Despite the importance of the nationality view, international statistics remain incomplete. In particular, comprehensive statistics on a nationality basis for portfolio investment and other positions are lacking. NBFIs manage vast portfolios through affiliates in financial centres, but identifying ultimate ownership or control requires piecing together various granular data sets (Coppola et al (2021); Beck et al (2024)). Closing these data gaps with official statistics would represent a major public good.

References

Aldasoro, I and T Ehlers (2017): "Risk transfers in international banking", *BIS Quarterly Review*, December.

Aldasoro, I, W Huang and E Kemp (2020): "Cross-border links between banks and non-bank financial institutions", *BIS Quarterly Review*, September.

Aldasoro, I, B Hardy and N Tarashev (2021): "Corporate debt: post-GFC through the pandemic", *BIS Quarterly Review*, June.

Avdjiev, S, M Chui and H S Shin (2014): "Non-financial corporations from emerging market economies and capital flows", *BIS Quarterly Review*, December.

Avdjiev, S, R McCauley and H S Shin (2016): "Breaking free of the triple coincidence in international finance", *Economic Policy*, vol 31, issue 87.

Barth, D, D Beltran, M Hoops, J Kahn, E Liu and M Perozek (2025): "The cross-border trail of the Treasury basis trade", *FEDS Notes*, October.

Beck, R, A Coppola, A Lewis, M Maggiori, M Schmitz and J Schreger (2024): "The geography of capital allocation in the euro area", *ECB Working Paper Series*, no 3007.

Bertaut, C, B Bressler and S Curcuru (2021): "Globalization and the reach of multinationals: implications for portfolio exposures, capital flows, and home bias", *Journal of Accounting and Finance*, vol 21, no 5.

Blanchard, O and J Acalin (2016): "What does measured FDI actually measure", *PIIE Policy Brief* 16-17.

Borio, C (2013): "The Great Financial Crisis: setting priorities for new statistics", *BIS Working Papers*, no 408.

Committee on the Global Financial System (CGFS) (2000): "Report of the Working Group on the BIS international banking statistics", *CGFS Papers*, no 15.

Coppola, A, M Maggiori, B Neiman and J Schreger (2021): "Redrawing the map of global capital flows: the role of cross-border financing and tax havens", *Quarterly Journal of Economics*, vol 136, no 3.

García Luna, P and B Hardy (2019): "Non-bank counterparties in international banking", *BIS Quarterly Review*, September

Gruić, B and P Wooldridge (2012): "Enhancements to the BIS debt securities statistics", *BIS Quarterly Review*, December.

Kaminsky, G and C Reinhart (1999): "The twin crises: The causes of banking and balance-of-payments problems", *American Economic Review*, vol 89, no 3.

Hardy, B, P McGuire and G von Peter (2024): "International finance through the lens of BIS statistics: bank exposures and country risk", *BIS Quarterly Review*, September.

Lane, P and G M Milesi-Ferretti (2018): "The external wealth of nations revisited: International financial integration in the aftermath of the global financial crisis", *IMF Economic Review* vol 66.

McCauley, R, P McGuire and P Wooldridge (2021): "Seven decades of international banking", *BIS Quarterly Review*, September.

McGuire, P and G von Peter (2009): "The US dollar shortage in global banking", *BIS Quarterly Review*, March.

McGuire, P, G von Peter and S Zhu (2024): "International finance through the lens of BIS statistics: residence vs nationality", *BIS Quarterly Review*, March.

Milesi-Ferretti, G M (2026): "The external wealth of nations database", the Brookings Institution, 6 February.

Pogliani, P, G von Peter and P Wooldridge (2022): "The outsize role of cross-border financial centres", *BIS Quarterly Review*, June.

Pogliani, P and P Wooldridge (2022): "Cross-border financial centres", *BIS Working Paper*, no 1035.

Box A

Financial centres and offshore activity

*Goetz von Peter and Philip Wooldridge*①

Offshore financial activity concentrates in financial centres, especially those that cater predominantly to non-residents. The provision of banking, fund management, business registry and other services to non-residents varies significantly across financial centres. These services give rise to external assets and liabilities, as recorded in residence-based statistics like the international investment position. The size of external positions relative to domestic economic activity thus provides a basis for distinguishing among financial centres and identifying those focused on offshore activity (Pogliani and Wooldridge (2022)).

Financial centres that mainly serve residents are *national centres*. They are usually home to domestic banks and the stock exchange, and they mainly intermediate domestic funding to foreign borrowers or foreign funding to domestic borrowers. The ratio of external intermediation to domestic activity is therefore low.

At the opposite extreme are *cross-border financial centres* (XFCs), which mainly serve non-resident counterparties and consequently have a very high share of international business (Graph A1, red dots). XFCs channel funds between economies, often via entities with a minimal physical presence, such as booking offices, special purpose vehicles and shell companies. They are neither the ultimate source nor final destination for investments and are usually embedded in small economies, as in the case of Bermuda and the Cayman Islands. While some equate XFCs with offshore centres, the latter have come to be associated with low taxes and light financial regulation, characteristics that are neither necessary nor sufficient to attract non-resident business (Pogliani et al (2022)).

Types of financial centres can be distinguished by the size of external positions[1] Graph A1

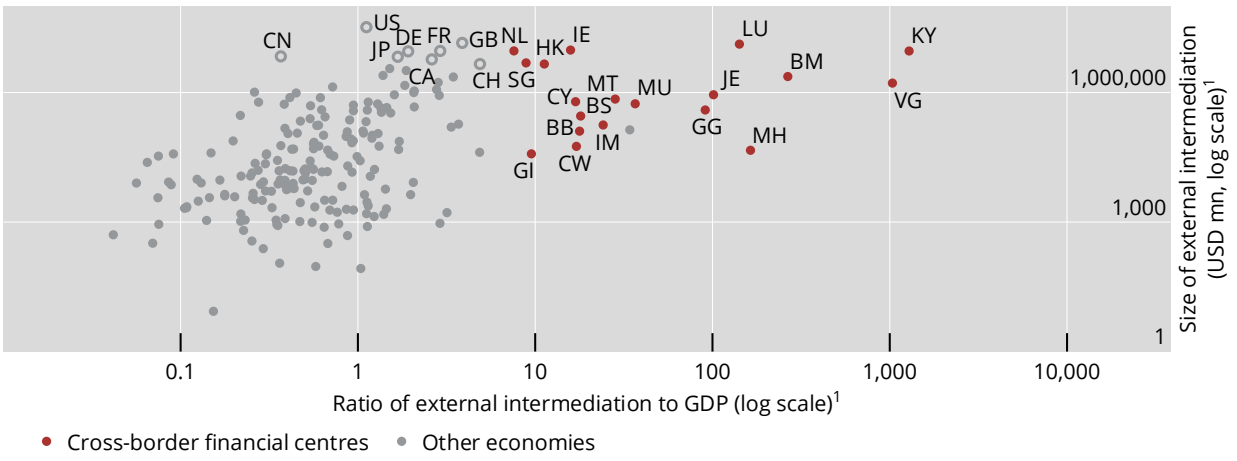

[1] At end-2024. External intermediation is measured as the minimum of an economy's external financial assets and liabilities.

Sources: Milesi-Ferretti (2026), based on Lane and Milesi-Ferretti (2018); Pogliani and Wooldridge (2022); authors' calculations.

Centres that combine the functions of national and cross-border centres at scale are commonly referred to as *global financial centres*. They are located in major economies, typically ones that issue reserve currencies, and are home to deep, liquid markets. Their international business is very large in absolute terms – indicated by economies towards the top of the y-axis in Graph A1 – but not necessarily relative to total economic activity – thus to the left of XFCs in Graph A1. London (GB) and New York (US) are the classic examples.

Global financial centres are key hubs for decisions about the international allocation of capital, but XFCs stand out as intermediaries of capital flows. From 1980 to 2023, XFCs' external positions grew from 12% of the global total to nearly 30% (Graph A2.A). At almost $60 trillion, their external positions are now 20 times larger than their combined GDP, which accounts for just 2.8% of global economic activity.

The prominence of cross-border financial centres (XFCs)

Graph A2

A. XFCs' rising share of external positions…[1]

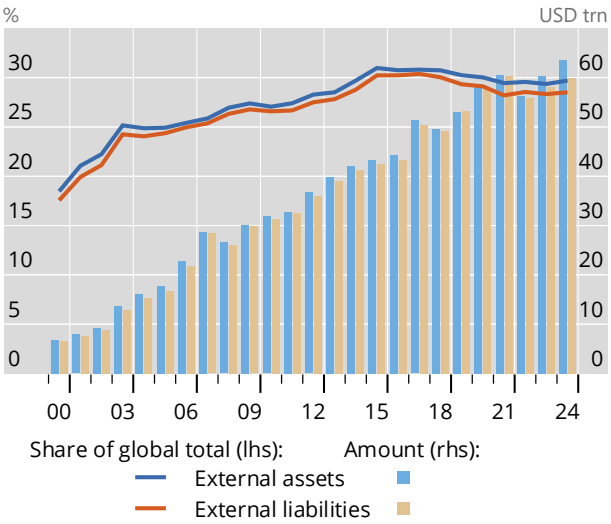

B. …furthered a divergence in the distributions of external positions and GDP[2]

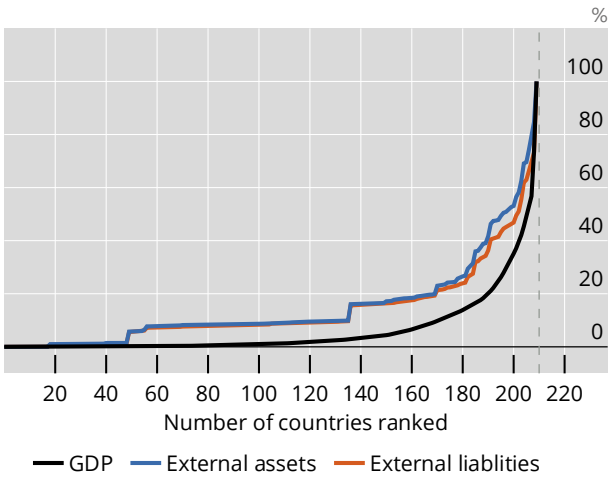

[1] The bars show the value of external assets and liabilities booked in cross-border financial centres, as defined in Graph A1. The lines show the corresponding shares of the global totals. [2] At end-2024. Countries are ranked in increasing order of size, from small to large GDP (x-axis); lines show cumulative shares in global totals. Each line ends on 100% of the global total (the sum over all countries).

Sources: Milesi-Ferretti (2026), based on Lane and Milesi-Ferretti (2018); authors' calculations.

Together with the growing heft of emerging market economies, the rise of XFCs has been a key driver of measured international financial integration (Lane and Milesi-Ferretti (2018)). Consequently, the accumulation of external positions in recent decades has proceeded with little relation to economic activity. Owing to activity in XFCs, the distribution of external positions worldwide has diverged markedly from the distribution of GDP (Graph A2.B).

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks.

Box B

Reallocation of liabilities to the ultimate obligor

*Iñaki Aldasoro, Bryan Hardy, Goetz von Peter and Philip Wooldridge*①

The nationality perspective highlights who controls the operations and balance sheet of a borrower, but the controlling parent is not necessarily the ultimate obligor responsible for repaying debt if the immediate borrower is unable to. Whether a borrower's liabilities are ultimately backed by its parent depends on corporate structures and guarantees. A parent is not obligated to support a subsidiary unless it has given an explicit guarantee (although they sometimes do even in the absence of a legally binding guarantee). Conversely, some firms guarantee the debt of unaffiliated firms (eg banks often do for trade finance). Thus, a comprehensive assessment of a country's financial vulnerabilities requires information beyond the borrower's nationality.

The BIS consolidated banking statistics are one of the few data sources that shed light on the reallocation of liabilities from immediate borrowers in one country to ultimate obligors in another. The statistics capture credit risk transfers used by banks, mainly guarantees but also collateral and credit derivatives (Aldasoro and Ehlers (2017)).② Some of these transfers, like credit derivatives, reduce banks' risk exposure but do not alter a borrower's debts; thus from a borrower's perspective, risk transfers overstate the reallocation of liabilities to ultimate obligors. Still, they can help to pinpoint where contingent liabilities may ultimately materialise.

Not surprisingly, a sizeable portion of offshore borrowing through affiliates in financial centres is guaranteed by entities elsewhere. For example, in late 2025 risk transfers reduced banks' claims on borrowers in the Cayman Islands by close to $300 billion and in Luxembourg by about $130 billion (Graph B1.A). Over time, outward risk transfers from the Cayman Islands have moved closely with inward risk transfers to the United States, suggesting that borrowing activity in the Cayman Islands was driven by US developments (Graph B1.B). As banks and firms from emerging market economies expanded internationally, their debt increased on an ultimate obligor basis. For example, since 2011 inward risk transfers to China and Korea have exceeded outward ones, indicating that guarantees given by Chinese and Korean firms and banks to their overseas affiliates have exceeded guarantees given by foreign firms to their own affiliates operating in China and Korea (Graph B1.B).

Borrowing in financial centres is often guaranteed by ultimate obligors elsewhere

Banks' net risk transfers,[1] in billions of US dollars — Graph B1

A. Net risk transfers, by counterparty country[2]

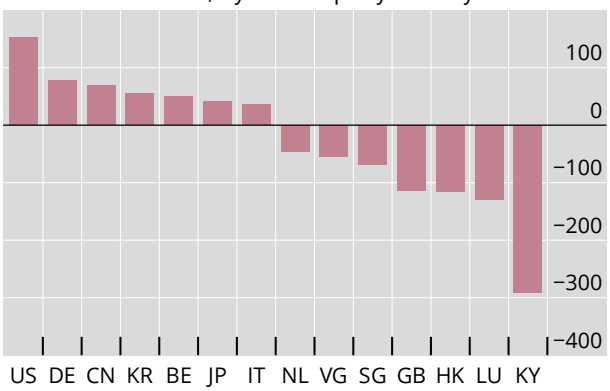

B. Evolution of net risk transfers over time

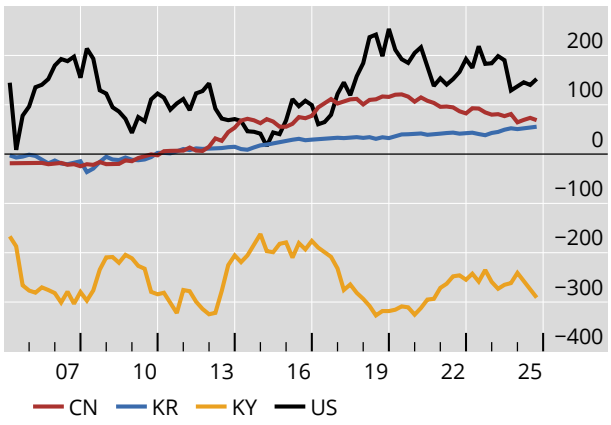

[1] Inward risk transfers that increase banks' claims on borrowers in the specified country minus outward risk transfers that reduce banks' claims on the country. [2] At end-September 2025, for selected counterparty countries.

Sources: BIS consolidated banking statistics; authors' calculations.

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks. ② Liabilities of banks' branches are regarded as being guaranteed even in the absence of an explicit guarantee.

Box C

Cross-border links between banks and NBFIs

*Iñaki Aldasoro and Bryan Hardy*①

The growing role of non-bank financial institutions (NBFIs) in intermediating cross-border flows has gone hand in hand with a proliferation of links between banks and NBFIs (Aldasoro et al (2020)) – links that often boost offshore activity. Owing to such links, banks' claims on NBFIs expanded much faster than overall cross-border bank lending, taking their share from 7% of overall lending in 2015 to 24% in 2025 (Graph C1.A, red line). This growth was largely concentrated on NBFIs located in the United States, the United Kingdom, the Cayman Islands and euro area financial centres (García Luna and Hardy (2019)).

NBFIs comprise a diverse set of market participants, ranging from investment funds and hedge funds to securitisation vehicles and central counterparties. The sectoral breakdown in the BIS international banking statistics is too coarse to distinguish among NBFIs, but the geographical distribution of bank-NBFI links hints at these NBFIs' underlying activities (Graph C1.B). Banks in the United States have strong links with Caribbean financial centres, especially the Cayman Islands, which probably reflects securitisation activity and prime brokerage (Barth et al (2025)). Banks in Japan also have sizeable claims on NBFIs in Caribbean financial centres, consistent with holdings of securities issued by securitisation vehicles (eg collateralised loan obligations). Central clearing probably contributes to strong ties between the United States and the United Kingdom, as both host large internationally active clearing houses for derivatives and repurchase agreements. Banks in the euro area and other advanced economies (eg Canada and Australia) have substantial claims on NBFIs in euro area financial centres (Ireland, Luxembourg and the Netherlands), in line with exposures to investment funds and other asset managers located there.

Banks, NBFIs and offshore activity

Graph C1

A. Cross-border bank claims on NBFIs continue to grow[1]

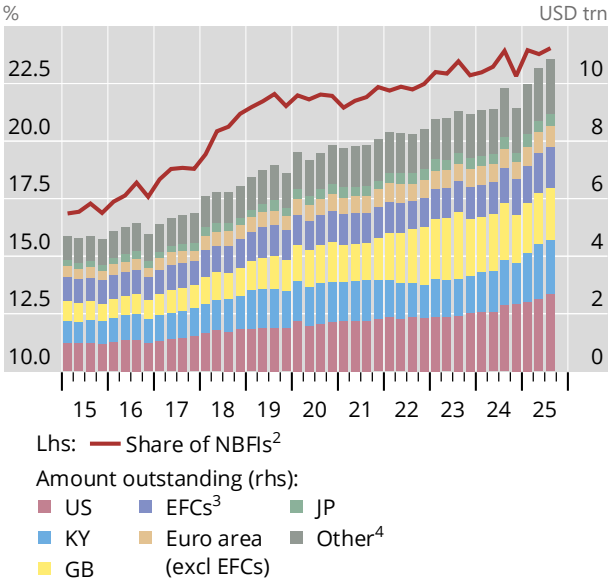

B. The network of cross-border bank claims on NBFIs[5]

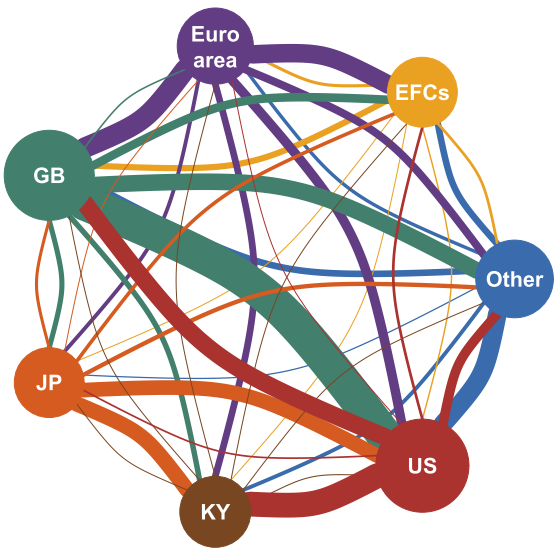

EFCs = euro area financial centres; NBFIs = non-bank financial institutions.

[1] Based on a varying number of reporting countries in respective quarters. [2] Share of NBFIs in the total counterparty sector. [3] IE, LU and NL. [4] Total excluding counterparties shown in the graph. [5] Country classification follows that in panel A. Nodes represent countries / country groups; size is proportional to the total value of cross-border positions as of end-Q3 2025; links connect bank to NBFI locations, with the colour indicating the bank's legal residence and the width reflecting the size of the claim.

Sources: BIS locational banking statistics by residence; authors' calculations.

① The views expressed here are those of the authors and not necessarily those of the BIS or its member central banks.

Abbreviations

Currencies

AED	United Arab Emirates dirham	MXN	Mexican peso
ALL	Albanian lek	MXV	Mexican unidad de inversión (UDI)
ARS	Argentine peso	MYR	Malaysian ringgit
AUD	Australian dollar	NAD	Namibian dollar
BHD	Bahraini dinar	NGN	Nigerian naira
BRL	Brazilian real	NOK	Norwegian krone
CAD	Canadian dollar	NZD	New Zealand dollar
CHF	Swiss franc	OTH	all other currencies
CLP	Chilean peso	PEN	Peruvian sol
CNY (RMB)	Chinese yuan (renminbi)	PHP	Philippine peso
COP	Colombian peso	PLN	Polish zloty
CZK	Czech koruna	RON	Romanian leu
DKK	Danish krone	RUB	Russian rouble
EUR	euro	SAR	Saudi riyal
GBP	pound sterling	SEK	Swedish krona
HKD	Hong Kong dollar	SGD	Singapore dollar
HUF	Hungarian forint	THB	Thai baht
IDR	Indonesian rupiah	TRY	Turkish lira
ILS	Israeli new shekel	TWD	New Taiwan dollar
INR	Indian rupee	USD	US dollar
ISK	Icelandic króna	VES	bolívar soberano
JPY	Japanese yen	VND	Vietnamese dong
KRW	Korean won	XOF	CFA franc (BCEAO)
MAD	Moroccan dirham	ZAR	South African rand

Countries

AE	United Arab Emirates	DE	Germany
AF	Afghanistan	DJ	Djibouti
AL	Albania	DK	Denmark
AM	Armenia	DM	Dominica
AO	Angola	DO	Dominican Republic
AR	Argentina	DZ	Algeria
AT	Austria	EA	euro area
AU	Australia	EC	Ecuador
AZ	Azerbaijan	EE	Estonia
BA	Bosnia and Herzegovina	EG	Egypt
BB	Barbados	ER	Eritrea
BD	Bangladesh	ES	Spain
BE	Belgium	ET	Ethiopia
BF	Burkina Faso	FI	Finland
BG	Bulgaria	FJ	Fiji
BH	Bahrain	FO	Faeroe Islands
BI	Burundi	FR	France
BJ	Benin	GA	Gabon
BM	Bermuda	GB	United Kingdom
BN	Brunei	GD	Grenada
BO	Bolivia	GE	Georgia
BR	Brazil	GG	Guernsey
BS	The Bahamas	GH	Ghana
BT	Bhutan	GI	Gibraltar
BY	Belarus	GN	Guinea
BZ	Belize	GQ	Equatorial Guinea
CA	Canada	GR	Greece
CD	Democratic Republic of the Congo	GT	Guatemala
CF	Central African Republic	GW	Guinea-Bissau
CG	Republic of Congo	GY	Guyana
CH	Switzerland	HK	Hong Kong SAR
CI	Côte d'Ivoire	HN	Honduras
CL	Chile	HR	Croatia
CM	Cameroon	HT	Haiti
CN	China	HU	Hungary
CO	Colombia	ID	Indonesia
CR	Costa Rica	IE	Ireland
CV	Cabo Verde	IL	Israel
CW	Curaçao	IM	Isle of Man
CY	Cyprus	IN	India
CZ	Czechia		

Countries (cont)

IQ	Iraq	MZ	Mozambique
IR	Iran	NA	Namibia
IS	Iceland	NC	New Caledonia
IT	Italy	NG	Nigeria
JE	Jersey	NL	Netherlands
JM	Jamaica	NO	Norway
JO	Jordan	NR	Nauru
JP	Japan	NZ	New Zealand
KE	Kenya	OM	Oman
KG	Kyrgyz Republic	PA	Panama
KH	Cambodia	PE	Peru
KR	Korea	PG	Papua New Guinea
KW	Kuwait	PH	Philippines
KY	Cayman Islands	PK	Pakistan
KZ	Kazakhstan	PL	Poland
LA	Laos	PT	Portugal
LB	Lebanon	PY	Paraguay
LC	St Lucia	QA	Qatar
LK	Sri Lanka	RO	Romania
LR	Liberia	RS	Serbia
LS	Lesotho	RU	Russia
LT	Lithuania	RW	Rwanda
LU	Luxembourg	SA	Saudi Arabia
LV	Latvia	SC	Seychelles
LY	Libya	SD	Sudan
MA	Morocco	SE	Sweden
MD	Moldova	SG	Singapore
ME	Montenegro	SK	Slovakia
MH	Marshall Islands	SI	Slovenia
MK	North Macedonia	SM	San Marino
ML	Mali	SR	Suriname
MM	Myanmar	SS	South Sudan
MN	Mongolia	ST	São Tomé and Príncipe
MO	Macao SAR	SV	El Salvador
MR	Mauritania	SZ	Eswatini
MT	Malta	TD	Chad
MU	Mauritius	TG	Togo
MV	Maldives	TH	Thailand
MW	Malawi	TJ	Tajikistan
MX	Mexico	TL	East Timor
MY	Malaysia	TM	Turkmenistan

Countries (cont)

TO	Tonga	UZ	Uzbekistan
TR	Türkiye	VC	St Vincent and the Grenadines
TT	Trinidad and Tobago	VE	Venezuela
TW	Chinese Taipei	VG	British Virgin Islands
TZ	Tanzania	VN	Vietnam
UA	Ukraine	ZA	South Africa
US	United States	ZM	Zambia
UY	Uruguay	1Z	British West Indies