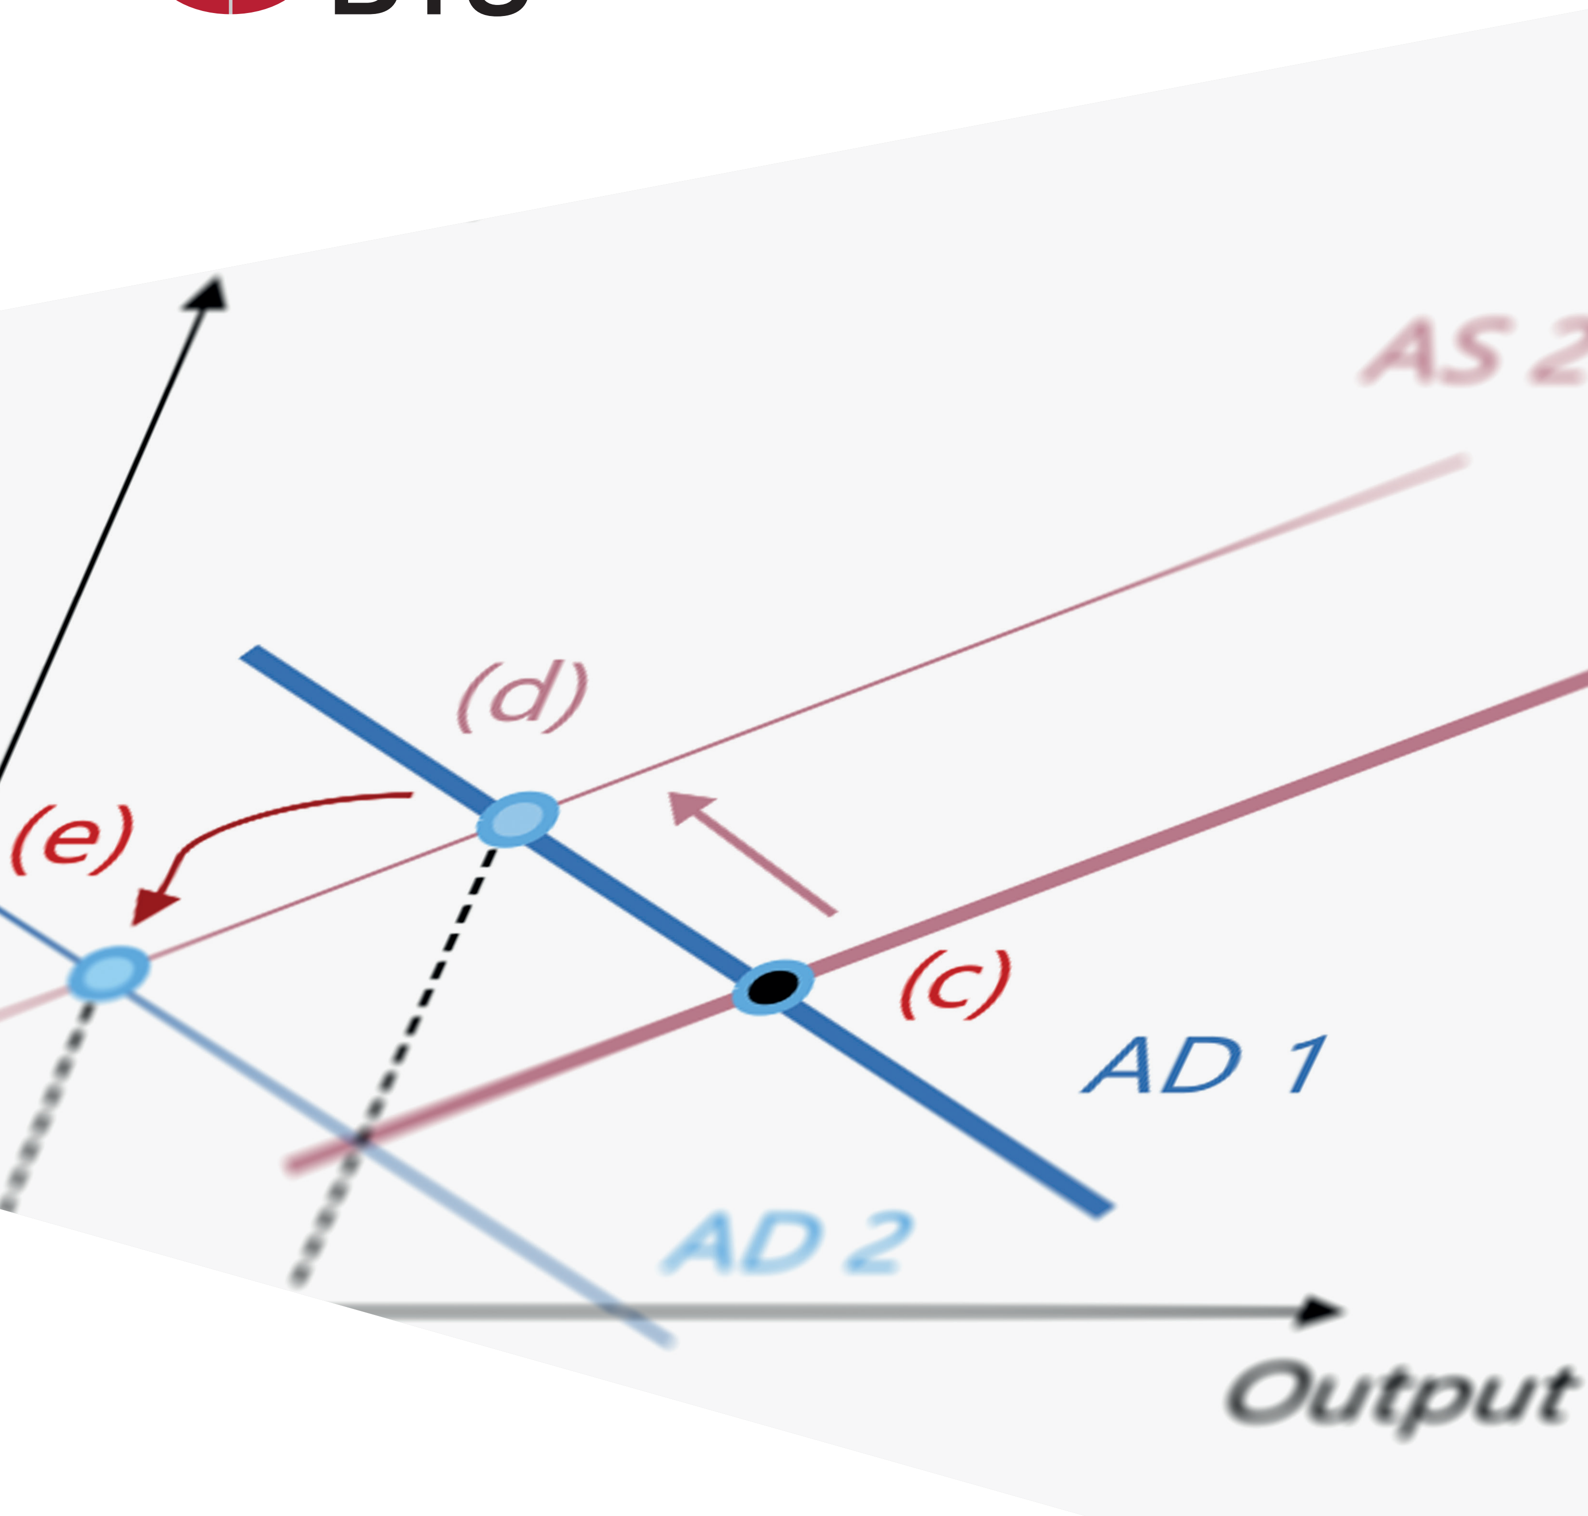


# BIS Quarterly Review

## International banking and financial market developments

December 2024

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_qt2412.htm).



ISSN 1683-0121 (print)
ISSN 1683-013X (online)

ISBN 978-92-9259-820-4 (print)
ISBN 978-92-9259-821-1 (online)

# BIS Quarterly Review

## December 2024

## International banking and financial market developments


Investor optimism prevails over uncertainty ........ 1

Bond markets diverge as US yields soar ........ 2

Box A: Growing uncertainty about terminal rates ........ 4

Equity and credit markets march on ........ 8

EME financial conditions tighten ........ 11

Box B: Negative interest rate swap spreads signal pressure in government debt absorption ........ 15

Technical annex ........ 18


## Special features


Targeted Taylor rules: monetary policy responses to demand- and supply-driven inflation ........ 19

*Boris Hofmann, Cristina Manea, Benoît Mojon*

Theory and central bank doctrine ........ 21

Empirical analysis ........ 24

Box A: Decomposing inflation into demand and supply components ........ 25

Targeted Taylor rules and the post-pandemic inflation surge ........ 29

Conclusions ........ 30

References ........ 32

Annex A ........ 34

Annex B ........ 35

Large language models: a primer for economists ........ 37

*Byeungchun Kwon, Taejin Park, Fernando Perez-Cruz and Phurichai Rungcharoenkitkul*

Introduction to large language models ........ 38

The central idea ........ 38

The techniques ........ 39

Box A: Transformer models ........ 40

Box B: In-context learning and prompting ........ 41

Putting LLMs to work ........ 41

Data organisation ... 42
Signal extraction ... 43
Quantitative analysis ... 44
Outcome evaluation ... 46
Illustration in a study of equity market drivers ... 46
Final considerations ... 50
References ... 51

Monetary policy and housing markets: insights using a novel measure of housing supply elasticity ... 53
*Ryan Banerjee, Denis Gorea, Deniz Igan and Gabor Pinter*
Measuring housing supply elasticity ... 55
Box A: Estimating the time-varying price elasticity of housing supply ... 56
Monetary policy transmission to house prices and rents ... 58
The influence of housing supply: evidence from a panel of countries ... 58
The influence of housing supply: potential channels ... 59
Policy considerations ... 61
References ... 62
Annex A: Rents and measurement of housing costs in CPIs ... 64
Technical annex ... 67

International finance through the lens of BIS statistics: the geography of banks' operations ... 69
*Bryan Hardy, Patrick McGuire, Goetz von Peter*
The architecture of global banking ... 71
The structure of banks' consolidated global operations ... 71
Box A: Geographic dimensions in BIS international banking statistics ... 72
Box B: Office types and the role of banks in facilitating capital flows ... 75
Implications of bank structure for policy ... 77
Bank resilience during the GFC ... 78
Internal markets: banks' intragroup funding ... 79
Conclusion ... 81
References ... 82
Annex A: The structure of banks' operations ... 84
Technical annex ... 85

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.

# 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 |
| BGN | Bulgarian lev | NGN | Nigerian naira |
| BHD | Bahraini dinar | NOK | Norwegian krone |
| BRL | Brazilian real | NZD | New Zealand dollar |
| CAD | Canadian dollar | OTH | All other currencies |
| CHF | Swiss franc | PEN | Peruvian sol |
| CLP | Chilean peso | PHP | Philippine peso |
| CNY (RMB) | Chinese yuan (renminbi) | PLN | Polish zloty |
| COP | Colombian peso | RON | Romanian leu |
| CZK | Czech koruna | RUB | Russian rouble |
| DKK | Danish krone | SAR | Saudi riyal |
| EUR | euro | SEK | Swedish krona |
| GBP | pound sterling | SGD | Singapore dollar |
| HKD | Hong Kong dollar | THB | Thai baht |
| HUF | Hungarian forint | TRY | Turkish lira |
| IDR | Indonesian rupiah | TWD | New Taiwan dollar |
| ILS | Israeli new shekel | USD | US dollar |
| INR | Indian rupee | VES | bolívar soberano |
| ISK | Icelandic króna | VND | Vietnamese dong |
| JPY | Japanese yen | XOF | CFA franc (BCEAO) |
| KRW | Korean won | ZAR | South African rand |
| MAD | Moroccan dirham | | |

## Countries

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

## Countries (cont)

| | | | |
|---|---|---|---|
| IN | India | MX | Mexico |
| IO | International organisations | MY | Malaysia |
| 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 |

## Countries (cont)

| | | | |
|---|---|---|---|
| TL | East Timor | UY | Uruguay |
| TM | Turkmenistan | UZ | Uzbekistan |
| TO | Tonga | VC | St Vincent and the Grenadines |
| TR | Türkiye | VE | Venezuela |
| TT | Trinidad and Tobago | VG | British Virgin Islands |
| TW | Chinese Taipei | VN | Vietnam |
| TZ | Tanzania | ZA | South Africa |
| UA | Ukraine | ZM | Zambia |
| US | United States | | |

# Investor optimism prevails over uncertainty

Despite lingering risks, investor optimism about the near-term outlook set the tone for financial markets during the review period.[1] The global economy seemed to be heading for a smooth landing, and the results of the US presidential election were conclusive. As a result, stock markets rose and credit spreads narrowed, easing global financial conditions. At the same time, rising government bond yields and an appreciating US dollar tightened them, pulling in different directions. Measures of risk premia and volatility ticked up in bond markets, amid signs that investors were pricing in higher fiscal and (geo-)political risks. However, the markets for risk assets mostly shrugged off these uncertainties and sentiment remained positive on balance.

Government bond yields generally rose, especially in the United States. The economy's strength there continued to surprise on the upside, pushing Treasury yields higher despite two consecutive policy rate cuts. More subdued activity in Europe meant that expectations of future rates were mostly unchanged and yields rose by less. Japanese government bond yields edged up as the Bank of Japan continued on its gradual path of policy normalisation. Rising US yields went along with a surge in the US dollar, a trend that intensified in the wake of the election.

Reemerging concerns about the fiscal situation in several jurisdictions, and quantitative tightening in others, added to the upward pressure on yields. Rising term premia, (more) negative swap spreads and widening sovereign spreads suggested that investors demanded a higher compensation to absorb additional debt supply.

Equity and credit markets posted substantial gains. US stocks reached all-time highs following the election, and Chinese stocks surged early in the review period in response to stimulus announcements. Elsewhere, stock market performance was more subdued. Corporate spreads continued to compress, and in some segments fell to multi-year lows. Buoyant risky asset markets counteracted the effect of rising yields and a strong US dollar on overall financial conditions. At the same time, there were signs that investors remained attuned to downside risks, with the costs of hedging in options markets pointing to lingering future uncertainties.

Financial conditions in emerging market economies (EMEs) tightened. With few exceptions, equity markets declined, currencies depreciated against the dollar and bond yields rose. China's stimulus announcement resulted in positive, albeit short-lived, spillovers to equity markets of EMEs with strong trade links to China. In currency markets, somewhat higher volatility reduced incentives for currency carry trades.

Slowing growth in China was reflected in depressed commodity prices. The impact was stronger in segments such as agricultural commodities and base metals, where demand from China is particularly large. Only the prices of gold and silver continued to rise, before retreating after the US election, possibly reflecting their perceived role as a hedge against geopolitical and inflation risks.

[1] The review period covers 7 September to 2 December 2024.

*Key takeaways*

- *Rapidly rising US yields and a soaring US dollar set the tone for fixed income and currency markets. Moves in other core bond markets followed a similar but more muted path, reflecting diverging macroeconomic conditions across regions.*
- *Risk-taking continued to be buoyant in equity and particularly credit markets, as investors largely shrugged off (geo-)political risks.*
- *EME financial conditions tightened, with higher bond yields, declining equity markets and headwinds posed by a stronger dollar.*

## Bond markets diverge as US yields soar

Over the review period, global fixed income markets resumed their divergence on the back of changing perceptions of the future paths of monetary policy and macroeconomic outlooks. While readings generally pointed at inflation converging towards targets, economic activity indicators painted a varied picture, as substantial strength in the United States contrasted with softness elsewhere. Additionally, potential concerns about the fiscal positions in some jurisdictions emerged as another key driver of bond markets. The US dollar appreciated initially with rising US yields and then surged further following the election.

Global bond markets diverge somewhat, with US yields in the lead[1] Graph 1

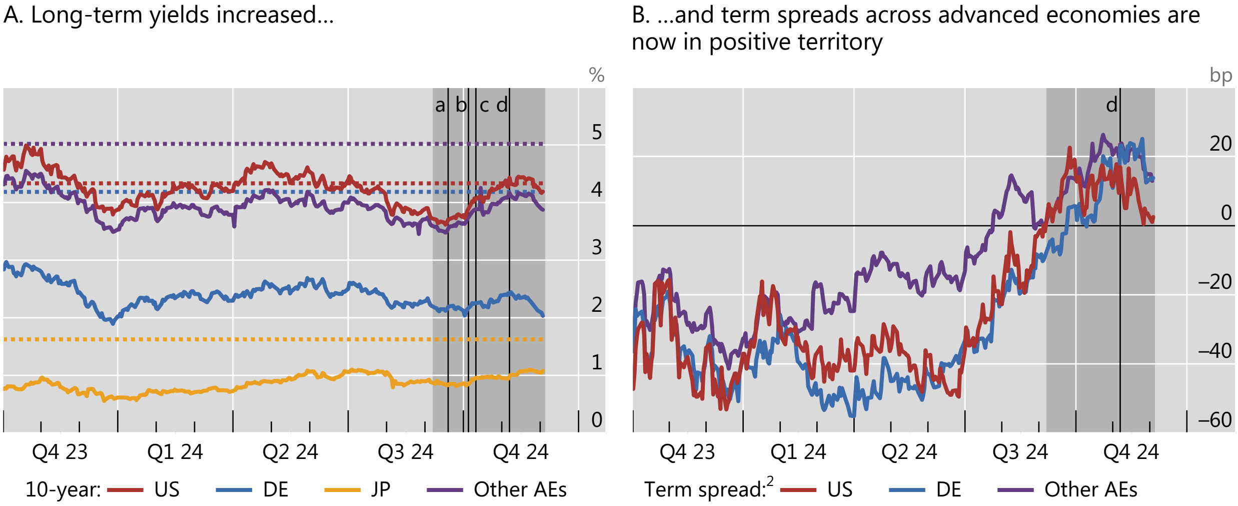


The shaded area indicates 7 September 2024–2 December 2024 (period under review). The dotted horizontal lines in panel A indicate the January 2007–June 2008 average.

[a] FOMC rate decision (18 September 2024). [b] US non-farm payrolls release (4 October 2024). [c] US CPI release (10 October 2024). [d] Day after the US presidential election (6 November 2024).

[1] Other advanced economies (AEs) based on simple average of AU, CA and GB. [2] Ten-year minus two-year.

Sources: Bloomberg; BIS.

Long-term yields rose across most advanced economies (AEs), led by soaring US Treasury yields for most of the review period. Despite the Federal Reserve's 50 basis point rate cut in September, US yields shot up in October. Ten-year Treasury yields rose by over 80 basis points from the trough reached on 16 September, before retracting somewhat late in the review period. Yields in other AEs followed a similar pattern, though with more muted movements overall, particularly in the euro area (Graph 1.A). The upward shift in yields initially went alongside a normalisation of the yield curve. Term spreads turned positive across AEs and yield curves broadly steepened, up until long-term yields retracted somewhat (Graph 1.B).

The macroeconomic backdrop, along with expected monetary policy paths, largely underpinned yield movements. Positive macroeconomic surprises continued to accumulate in the United States, but were more short-lived elsewhere (Graph 2.A). As data releases highlighted a resilient labour market and somewhat persistent inflation (Graph 2.B), market participants lifted the expected 2025 path of the federal funds rate (Graph 2.C). Further upward revisions took place following the US presidential election as markets digested the policy changes which were likely to ensue. By the end of the review period, the US term structure flattened again, as near-term yields rose on expectations of a strong economy and fewer policy rate cuts next year. These revisions coincided with signs of increased uncertainty about the level of terminal rates (Box A).

Revisions to the expected monetary policy paths and resulting yield movements were not as pronounced elsewhere. The macroeconomic backdrop in other major AEs was less benign. With weaker growth expected in the euro area in particular, and with inflation even falling below target in September, investors priced in larger rate cuts by the ECB. In Japan, the central bank signalled that it would continue to normalise policy very gradually, as inflation stayed somewhat above target. Long-term Japanese government bond yields rose modestly during the review period, reflecting gradual policy normalisation by the Bank of Japan.

Economic developments and policy expectations largely drive changes in yields — Graph 2

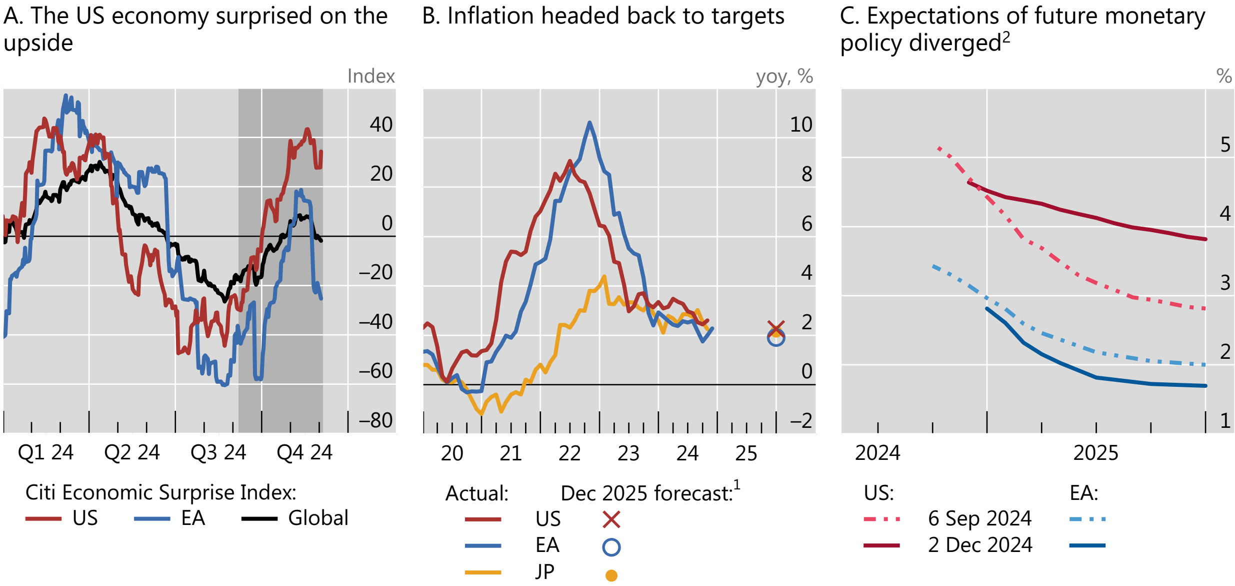


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[1] November 2024 forecast for December 2025. [2] Futures curve of policy rates.

Sources: Bloomberg; Consensus Economics; Macrobond; national data; BIS.

Box A

## Growing uncertainty about terminal rates

*Matteo Aquilina, Marco Lombardi and Sonya Zhu*①

With inflation on track towards targets and global economic activity remaining strong, central banks began easing the policy stance earlier this year. At the same time, the uncertainty over the level at which policy rates will eventually settle has been growing. This box documents how disagreement over "terminal rates" has recently evolved, based on the views of policymakers and professional forecasters.

Market participants have held a wide range of views about the short- to medium-term monetary policy outlook since 2021.② At first, as central banks embarked on policy rate hikes, participants wondered how much tightening would be necessary to rein in inflation. Subsequently, as inflation started declining steadily, market commentary revolved around the question of how long rates would be kept high to ensure plain sailing to targets. More recently, as the easing phase began, the debate in markets has shifted towards terminal rates.

The terminal rate can be viewed as the level of the policy rate that is consistent with inflation at target and economic activity at potential. As such, it is simply the nominal equivalent of the so-called natural rate of interest, or r*.③ In this sense, steering policy towards the terminal rate means aiming for a neutral policy stance, ie one that is neither tight nor loose. Similarly to r*, the terminal rate cannot be observed directly and must be inferred from the dynamic interplay of the variables that help define it: policy rates, inflation and economic activity.

### Views of terminal rates among forecasters and policymakers are diverging

Graph A1

A. Forward rates | B. Consensus forecasters disagree more[1] | C. Long-term policy rates in the US

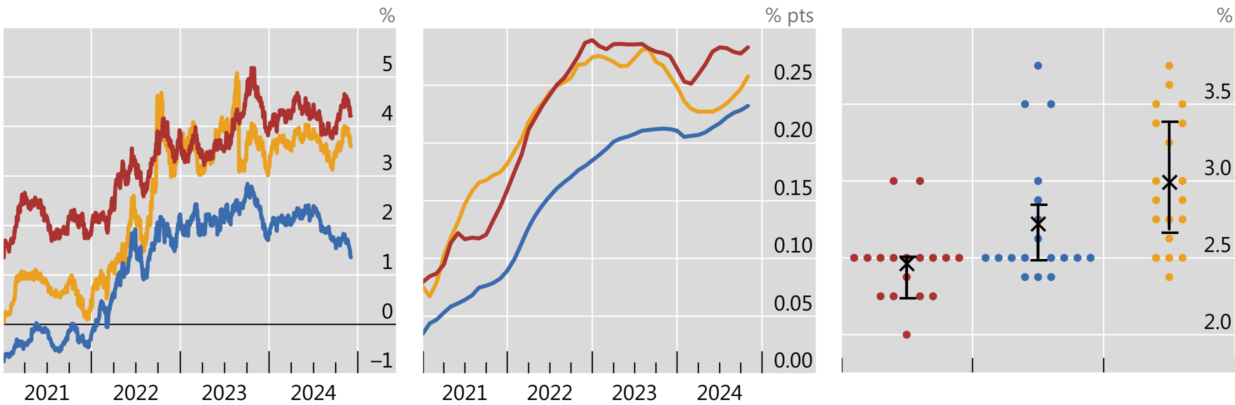


Three-month rate in five years: USD, EUR, GBP

Terminal rate estimates (standard deviation): US, EA, GB

Fed funds rate target range or level, longer run: Dec 2021, Dec 2023, Sep 2024; Interquartile range; Mean

[1] For further details on the calculations, see Lombardi, Schrimpf and Zhu (forthcoming).

Sources: Lombardi, Schrimpf and Zhu, "Heterogenous beliefs and the central bank reaction function", *BIS Working Papers*, forthcoming; Board of Governors of the Federal Reserve System; Bloomberg; Consensus Economics; authors' calculations.

The global economy's apparent resilience despite the sharp policy tightening has led to upward revisions in proxies for terminal rates. Even with the caveat that they also reflect term premia, forward rates give a broad sense of bond market participants' view of how future short rates will evolve. A reasonable proxy for the terminal rate is the short-term rate expected to prevail five years ahead. Over the past few months, such short-term forward rates rose steadily, back to the levels of late 2023 (Graph A1.A). This has been hinting at a shallower easing trajectory and at an overall higher level of rates at the end of the loosening phase.

In parallel with an upward shift in expectations of terminal rates, the uncertainty surrounding them also increased. One way to gauge this is to consider professional forecasters' views on the policy stance across different jurisdictions. Professional forecasters in surveys are typically not directly asked about their views on policy rates in the long run, but one can use their expectations on interest rates, inflation and output to gauge their perceptions of the monetary policy rule, including the perceived level of neutral nominal rates. This can be done by stripping their short-run monetary policy expectations from the component that they ascribe to the cyclical position of the economy, ie their short-run expectations on inflation and output.④ On this basis, the estimated standard deviation of professional forecasters' perceived terminal rates rose across different jurisdictions from 2021 to 2022 in tandem with inflation. It then started declining somewhat as inflation receded, but has picked up again since mid-2024 in the euro area, the United Kingdom and the United States, even with inflation well on track towards targets (Graph A1.B).

The growing divergence of professional forecasters' views on terminal rates reflects uncertainty about a variety of factors. As argued above, one is the somewhat surprising resilience of economic activity to higher rates. Another relates to the outlook of inflation and to the central banks' expected response to possible future over- or undershooting of targets. But disagreement can also reflect higher uncertainty about the speed and strength of the transmission of monetary policy to economic activity.

The growing uncertainty among professional forecasters also echoes policymakers' own diverging views on terminal rates. This is evident in the United States, where the future policy rate forecasts of Federal Open Market Committee members are regularly disclosed through the Fed dot plot: members' views of the level of policy rates in the *long term* are a reasonable proxy for expected terminal rates. The upward revisions in those rates that took place as the economy repeatedly surprised on the upside coincided with a growing dispersion. At the end of 2023, expected terminal rates were still mostly clustered around 2.5%, as they had been two years earlier. By September 2024, the range had become substantially wider (Graph A1.C).

The growing divergence in views on terminal rates among both policymakers and professional forecasters underscores the challenges central banks face in navigating the current economic waters. Given the associated uncertainties, it also highlights the importance of relying on robust monetary policy frameworks.⑤

① The views expressed are those of the authors and do not necessarily reflect the views of the BIS. ② See M Aquilina, M Lombardi and S Zhu, "The return of monetary policy uncertainty", *BIS Quarterly Review*, March 2024. ③ For an introductory discussion, see G Benigno, B Hofmann, G Nuño and D Sandri, "Quo vadis, r*? The natural rate of interest after the pandemic", *BIS Quarterly Review*, March 2024. ④ For further details, see M Lombardi, A Schrimpf and S Zhu, "Heterogeneous beliefs and the central bank reaction function", *BIS Working Papers*, forthcoming. ⑤ For an overview of the issues, see BIS, "Monetary policy in the 21st century: lessons learned and challenges ahead", *Annual Economic Report 2024*, Chapter II, 2024.

The US dollar soared and the yen touched lows not seen since late July, but developments in fixed income markets only partly explained these movements. The dollar appreciated with rising US yields early in the review period, but then surged even further on the US election outcome (Graph 3.A), even as US long-term yields retracted somewhat. Among AE currencies, the yen depreciated notably vis-à-vis the US dollar (Graph 3.B), by some measures in excess of what would be expected based on the yield gap in the respective bond markets. Net short positions in yen currency futures started rebuilding (Graph 3.C), suggesting investor positioning for further yen weakness. Typically, short positions by financial investors, especially hedge funds, suggest a rise in yen carry trades. However, this time leveraged funds held a smaller share of yen shorts due to increased volatility making such bets riskier (see below). Thus, the current positioning may primarily reflect greater currency hedging by asset managers holding yen assets.

Some indicators in fixed income markets pointed to greater investor uneasiness. For one, the implied volatility of US Treasuries reached the highest level of the year (Graph 4.A) before falling sharply following the election. At the long end of the yield curve, estimates of the term premium edged up through the review period, indicating that investors demanded greater risk compensation for holding long-term US government debt (Graph 4.B).[2]

[2] Model-based estimates indicate that the increase in the term premium was not due to inflation risk, which has diminished, but to an increase in the real premium.

## US dollar soars and yen depreciates to levels not seen since July

Graph 3

A. US dollar appreciated with rising yields and US election outcome

B. Yen depreciated more than warranted by the yield differential

C. Net short positions in yen futures started rebuilding

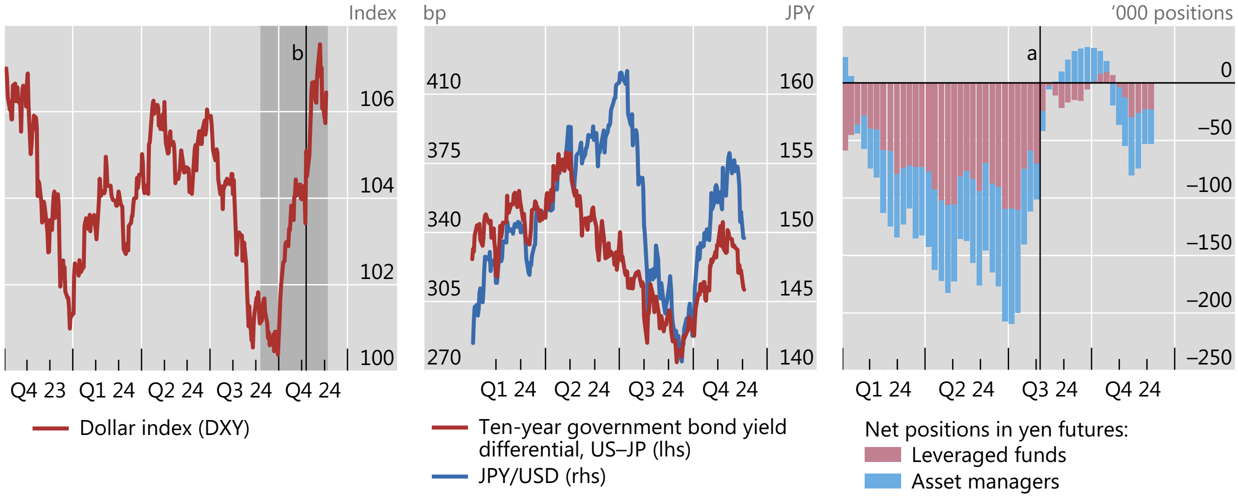


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[a] Turbulence in the Japanese equity market (5 August 2024). [b] Day after the US presidential election (6 November 2024).

Sources: Bloomberg; LSEG Datastream; BIS.

Higher measures of risk in fixed income markets coincided with tell-tale signs of a possible supply glut of government bonds. Interest rate swap spreads – the swap rate minus the government bond yield – fell rapidly, indicating that government bonds had become relatively cheaper (and their yields relatively higher). The phenomenon was widespread across currencies and maturities, with the euro and Japanese yen spreads joining their US counterparts in negative territory, even at those shorter maturities where spreads had previously been positive (Graph 4.C). Negative

## Cautionary signs in fixed income markets

Graph 4

A. US Treasury volatility increased

B. The US term premium spiked

C. Negative swap spreads prevailed

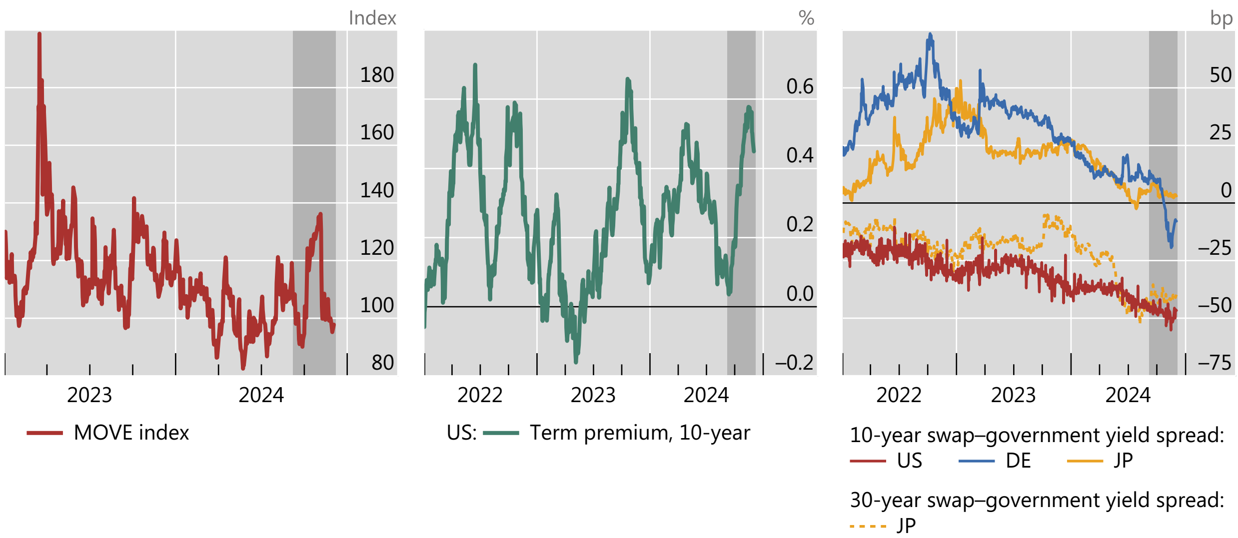


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

Sources: D Kim and J Wright "An arbitrage-free three-factor term structure model and the recent behaviour of long-term yields and distant-horizon forward rates", *FEDS Working Papers*, no 2005-33, 2005; Board of Governors of the Federal Reserve System; Bloomberg; BIS.

swap spreads appeared to reflect pressures on investors and intermediaries due to the need to absorb more government debt supply in the near future (Box B).

Conditions in short-term funding markets over the review period reflected ongoing quantitative tightening and an increasing shift towards a relative abundance of collateral. In the United States, repo spreads – defined here as the difference between the rate on an overnight repo and the effective federal funds rate – spiked on a few occasions, most notably in September (Graph 5.A). This seemingly reflected the greater repo financing needed to absorb large Treasury issuance as well as constraints on dealers' balance sheets linked to end-of-quarter regulatory reporting pressures. Repo spreads in Japan, while still negative, showed signs of normalisation in recent months (Graph 5.B), pointing to government collateral becoming less scarce, despite still very ample liquidity. In the euro area too, the reversal of collateral scarcity drove repo rates higher.[3] Large sovereign issuance and a dwindling ECB footprint contributed to the gradual increase in repo rates over the past 18 months. This shift was also visible in the rapid reduction in take-up of the ECB Securities Lending Programme, particularly against cash collateral (Graph 5.C): investors had less need to turn to the central bank as they were able to obtain collateral from other sources.

Several indicators pointed to the re-emergence of fiscal concerns. US sovereign credit default swap (CDS) spreads widened in advance of the November presidential and congressional elections, amid worries about the fiscal implications of the uncertain outcomes (Graph 6.A). The subsequent narrowing may have reflected a drop in the likelihood of another debt ceiling impasse, as the next president's party gained a congressional majority. In the euro area, French sovereign spreads remained elevated (Graph 6.B), as the newly appointed government announced larger than

Money markets shift towards collateral abundance

Graph 5

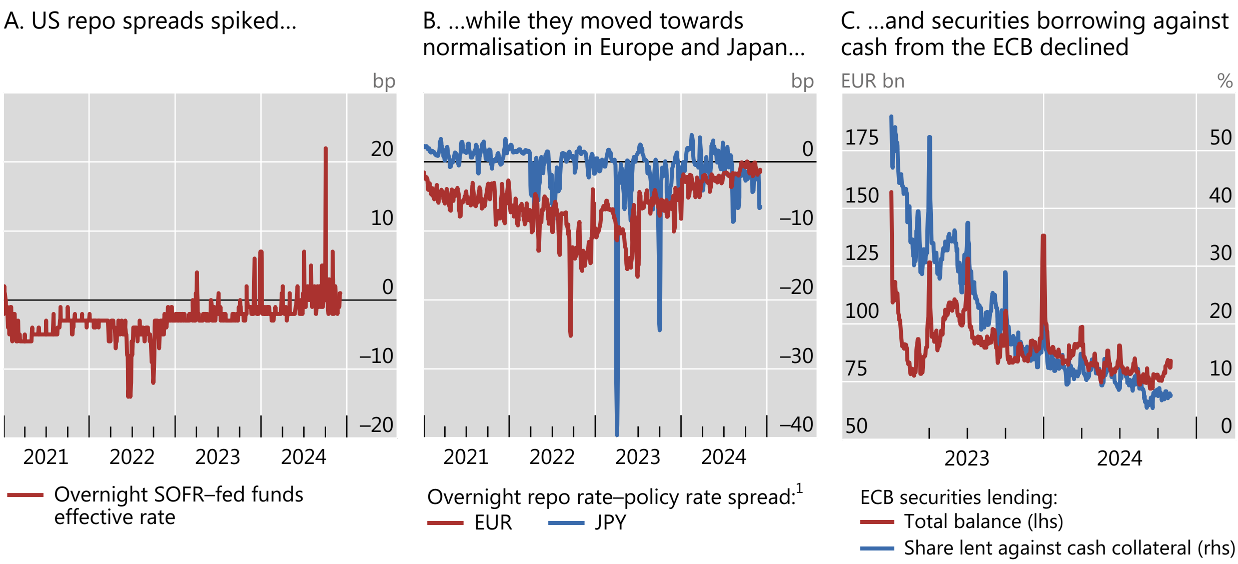


SOFR = secured overnight financing rate.

[1] Five-day moving average excluding quarter-ends.

Sources: ECB; Bloomberg; LSEG Datastream; national data; BIS.

[3] The abundance of collateral is reflected in higher repo rates as lenders demand a higher rate of interest to lend cash against collateral, such as government bonds.

Fiscal concerns become a greater focal point

In basis points

Graph 6

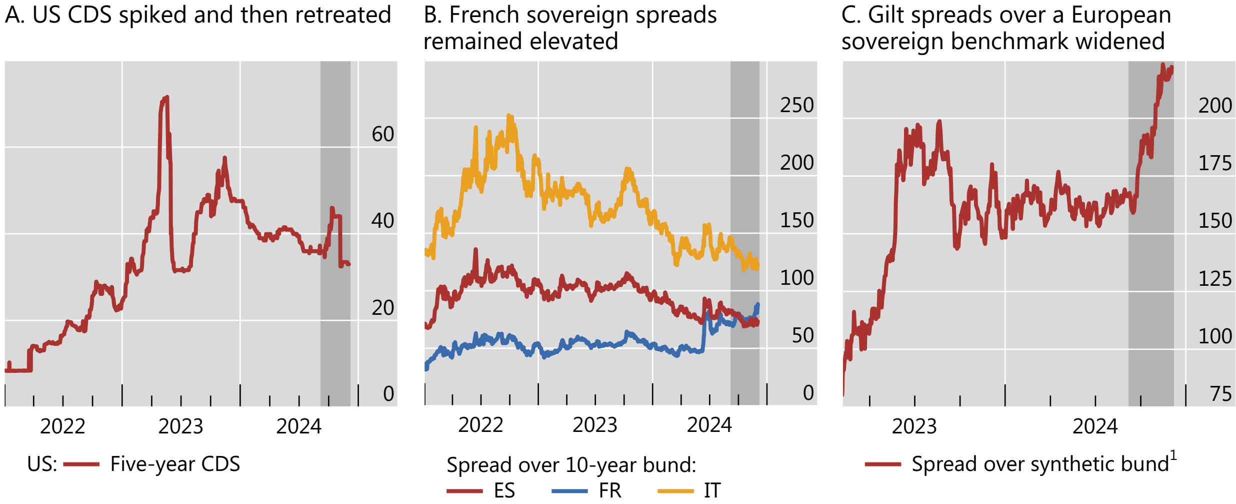


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[1] Difference between the yield on 10-year UK government bonds and the yield on 10-year German government bonds adjusted for currency differences using a EUR/GBP 10-year currency swap.

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

expected fiscal deficits, while its ability to command a majority in parliament was itself in doubt. The fiscal outlook also appeared expansionary in Japan on the back of the October parliamentary election results. And it became so in the United Kingdom, where the government announced a relaxation of fiscal rules and more borrowing. This sent yields on gilts higher relative to a synthetic benchmark based on German bunds (expressed in the same currency by using FX swaps to convert the yield in euros into a gilt-equivalent yield) (Graph 6.C).

## Equity and credit markets march on

Equity and credit markets were undeterred by political and fiscal risks, posting gains as the global economy continued to show strength. Corporate spreads compressed further, even in the euro area and in Asian markets, where the growth outlook was more mixed. The credit spread compression counteracted the effect of rising government bond yields and a stronger US dollar on overall financial conditions. At the same time, investors appeared to be more attuned to downside risks, at least in equity markets, with the VIX lingering at higher levels following the August turbulence, before retreating somewhat after the US election.

Stock market performance varied significantly across major economies. US equities recorded fresh highs (Graph 7.A), as market valuations were buoyed by the Federal Reserve's 50 basis point September rate cut and by the economy's strength. Stocks, especially small caps, shot up further on 6 November once the outcome of the US election became clear. Bitcoin followed a similar pattern, reaching all-time highs. Across the Atlantic, the more subdued performance of major European equity markets was in line with the relatively weaker economy. Expected earnings for European-listed companies deteriorated further and trailed those of their US peers

Some stock markets reached new heights despite temporarily higher volatility — Graph 7

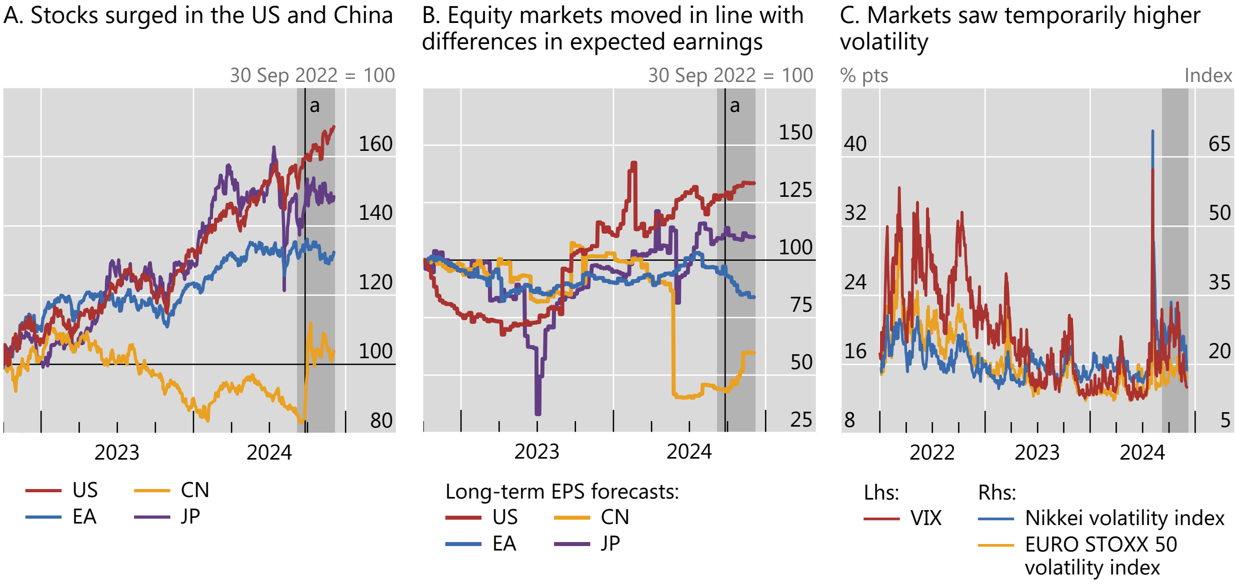


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[a] China stimulus announcement (24 September 2024).

Sources: Bloomberg; LSEG Datastream; BIS.

by a significant margin (Graph 7.B). To some extent, their relative weakness also reflected a greater weight of more cyclical industrial and energy stocks in the index.

In China, equities rose substantially following the announcement of a large policy stimulus package in late September. That said, the impact on earnings expectations was more muted, as authorities released information on the content of the package in batches. Following the initial jump, Chinese equities moved mostly sideways.

Notwithstanding the large gains of US equity indices, investors appeared to be growing warier of potential risks. After the volatility of early August subsided, the VIX continued to fluctuate within higher ranges (Graph 7.C). It then retreated as the (near-term) uncertainty surrounding the US election was resolved. However, the term structure of the VIX became upward sloping again, with higher expected medium-term volatility.

Conditions in credit markets around the globe remained unusually accommodative. Corporate credit spreads stayed compressed relative to their historical watermarks, in both the investment grade and high-yield segments (Graphs 8.A and 8.B), in some cases reaching lows not seen since the mid-2000s. Indeed, some analysts started referring to conditions in corporate funding markets as the "valuation conundrum".

Several underlying forces appeared to put a lid on corporate spreads. For one, debt default rates moderated in the United States (Graph 8.C). Furthermore, corporates' immediate refinancing needs subsided as the maturity wall was pushed out until late 2025–26. In addition, US corporate debt markets continued to see strong foreign demand, with almost two years of consecutive net foreign purchases, according to Treasury International Capital data. Finally, compressed spreads also reflected higher government bond yields rather than lower corporate yields.

## Credit markets pull ahead, with spreads at multi-year lows[1] Graph 8

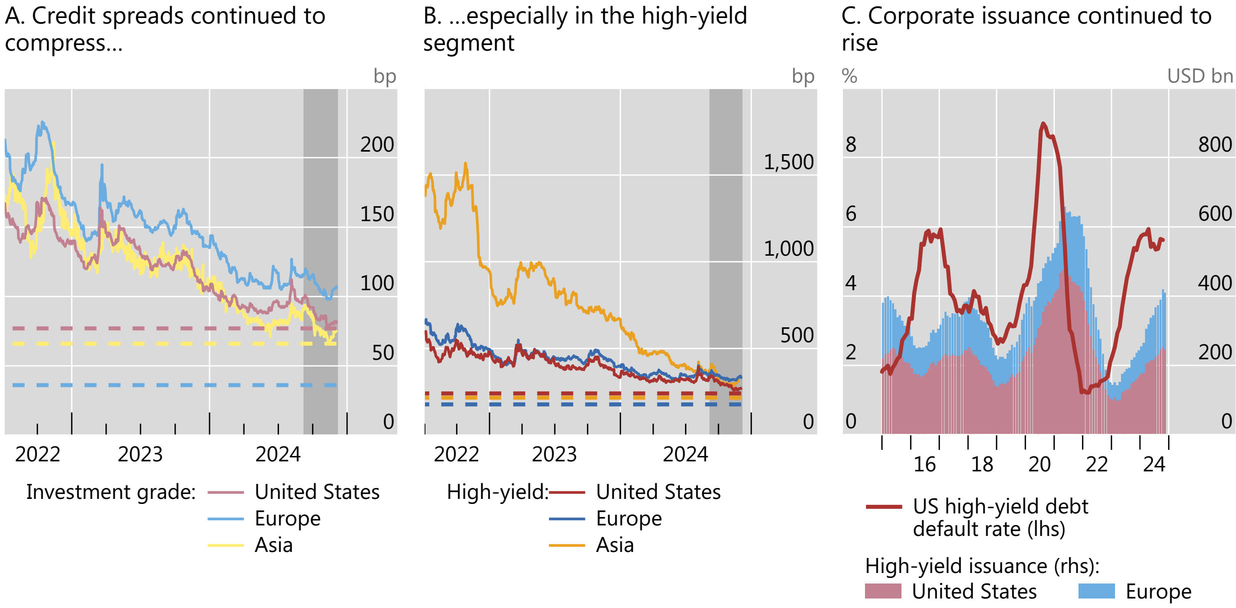


The shaded area indicates 7 September 2024–2 December 2024 (period under review). The horizontal dashed lines in panels A and B indicate 2005–current minima.

[1] See technical annex for details.

Sources: Dealogic; ICE Data Indices; Moody's; BIS.

Bank lending conditions eased amid a surge in bank equity prices in the United States. Stock prices of banks began to rally following the September policy rate cut, especially for regional banks with large commercial real estate (CRE) portfolios. The outperformance of CRE-exposed banks reflected shifting investor sentiment towards

## Bank stocks rally and lending terms ease Graph 9

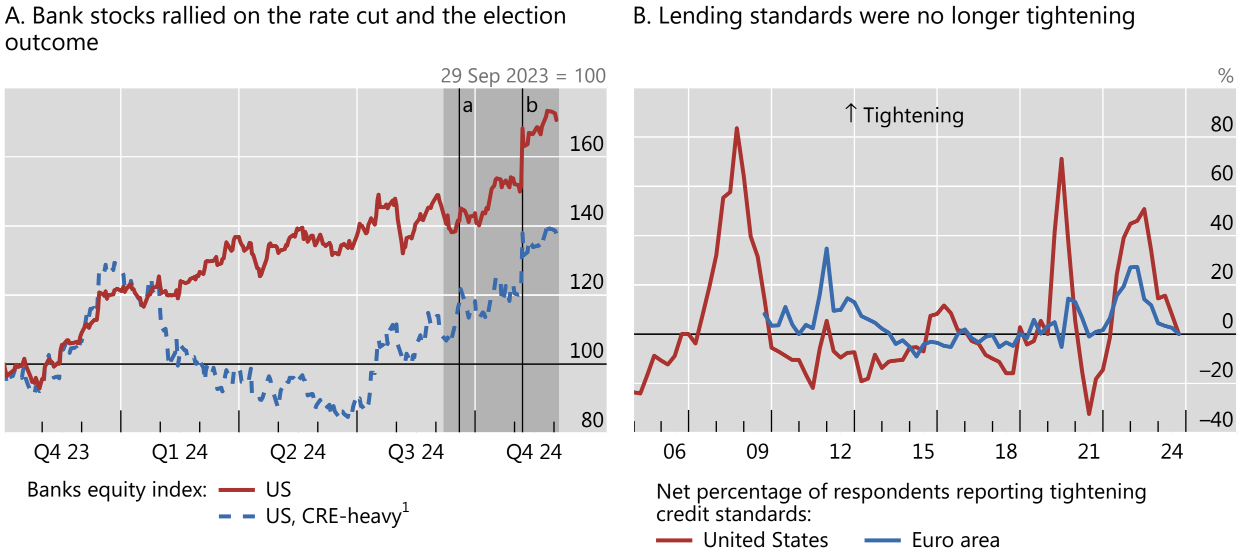


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[a] FOMC rate decision (18 September 2024). [b] Day after the US presidential election (6 November 2024).

[1] CRE = commercial real estate. See technical annex for details.

Sources: Board of Governors of the Federal Reserve System; LSEG Datastream; BIS.

risk in the sector (Graph 9.A). Bank stock prices rose further after the US presidential election, as investors looked to some regulatory easing. Rising stock prices went hand in hand with easing lending standards (Graph 9.B); declining non-performing loan ratios supported this shift in both the United States and Europe.

## EME financial conditions tighten

EMEs faced headwinds and financial conditions tightened. With few exceptions, equity markets declined, currencies depreciated and bond yields rose. China's stimulus announcement caused visible, albeit short-lived, positive spillovers to equity markets of EMEs with stronger trade links to China. Some EME currencies also depreciated on domestic political and fiscal uncertainties.

Bond yields in EMEs diverged across regions, reflecting differences in the macro backdrop, monetary policy cycles and exposure to US developments. Yields in major Latin American economies increased notably (Graph 10.A), tracking their US counterparts more closely than those in other EMEs. Bond yields in Mexico proved especially sensitive to US developments, while those in Brazil reflected mainly strong domestic economic activity and monetary tightening amid rising inflationary pressures. Most Latin American countries also maintained an expansionary fiscal stance, putting further upward pressure on yields. In emerging Asia, by contrast, many central banks cut policy rates for the first time in the cycle and yields rose substantially less. At the end of the spectrum, Chinese government bond yields stayed low, amid subdued growth and policy easing.

EME bond yields drift higher and currencies depreciate — Graph 10

A. Bond yields mostly increased — B. Exchange rates depreciated as US yields rose — C. Implied volatility rose

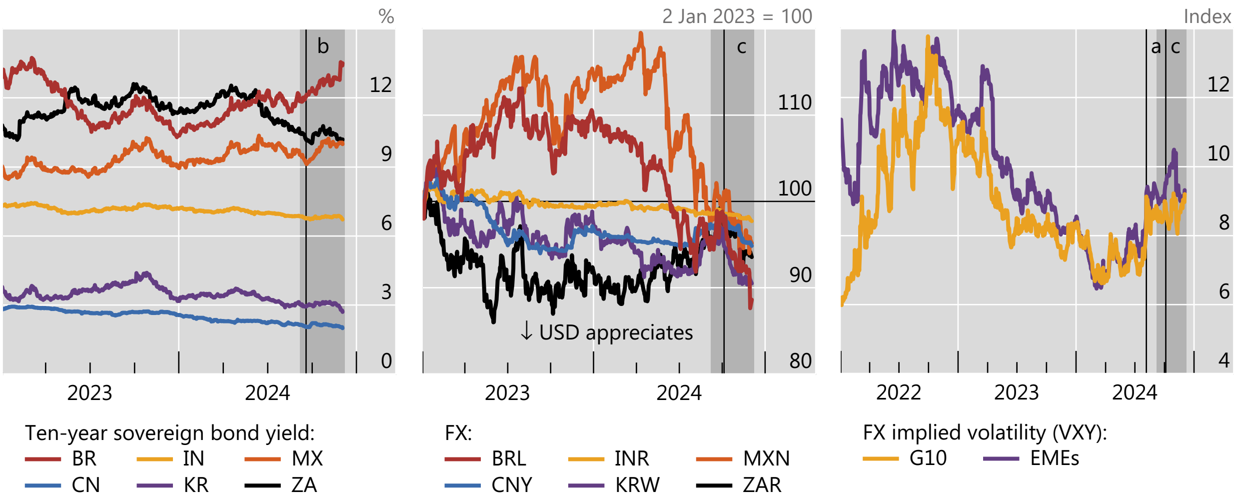


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[a] Japanese equity market turbulence (5 August 2024). [b] FOMC rate decision (18 September 2024). [c] US non-farm payrolls release (4 October 2024).

Sources: Bloomberg; JPMorgan Chase; national data; BIS.

EME currencies generally depreciated against the greenback, while their volatility increased. The depreciation was particularly large for Latin American currencies. The Mexican peso depreciated further on the likelihood of trade tariffs under the new US administration (Graph 10.B). The Brazilian real tumbled in December as the highly anticipated measures to curb public spending were thrown into doubt. In Asia-Pacific, the Korean won was briefly shaken by a bout of domestic political instability. Still, while EME currency volatility picked up in October (Graph 10.C), it decreased somewhat for most EME currencies with the US election results, in line with retreating near-term implied volatility in other markets.

Higher volatility in foreign exchange markets during the review period reduced the incentive for carry traders to rebuild their positions following the August unwind. Carry-to-risk ratios declined substantially, and mostly because of the increase in volatility (Graph 11.A). Furthermore, option-implied measures of skewness showed that investors have grown warier of the potential downside risks associated with these strategies: jurisdictions with the highest interest rate differential to the United States are those where the currency depreciation risk is higher (Graph 11.B).

EMEs with strong links to China saw their equity markets benefit from the Chinese government's stimulus announcement. Following China's stock market jump, which erased more than a year of losses in a single day (Graph 12.A), other markets quickly followed suit. The impact was especially noticeable in Brazil and some other Latin American commodity exporters such as Chile. Within emerging Asia, the Korean and Thai stock markets were among the most affected. However, such spillovers were short-lived, and EME equity indices declined rapidly thereafter.

## Higher volatility reduces incentives for currency carry trades

Graph 11

A. Carry trades less attractive due to higher volatility[1]

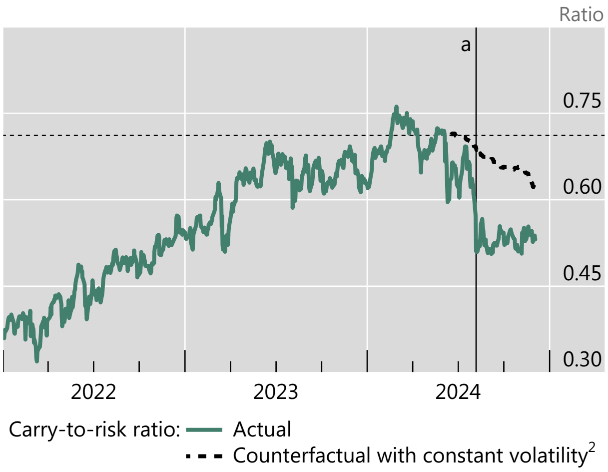


B. Currency crash risk in options explained the carry

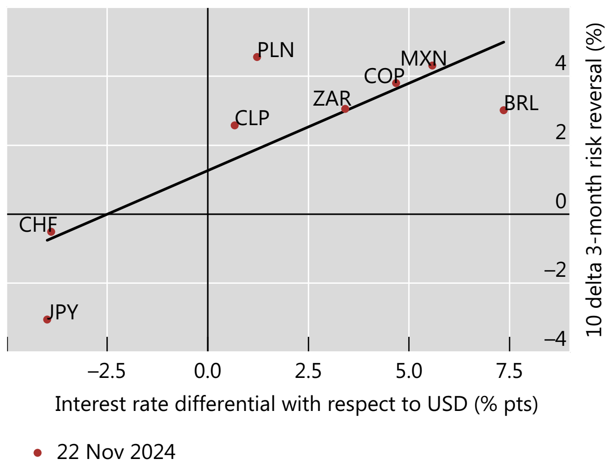


[a] Turbulence in the Japanese equity market (5 August 2024).

[1] Implied volatility of three-month at-the-money JPY currency options. Median across: BRL, MXN, PLN and ZAR. [2] Implied volatility fixed at the level on 3 June 2024. The dashed horizontal line denotes the value on 3 June 2024.

Sources: Bloomberg; JPMorgan Chase; LSEG Datastream; Macrobond; BIS.

EME stock markets other than China generally decline amid volatile capital flows — Graph 12

A. Some equity markets reacted to China's stimulus announcement

B. China saw a spike in portfolio equity inflows

C. Other EMEs mostly saw bond inflows

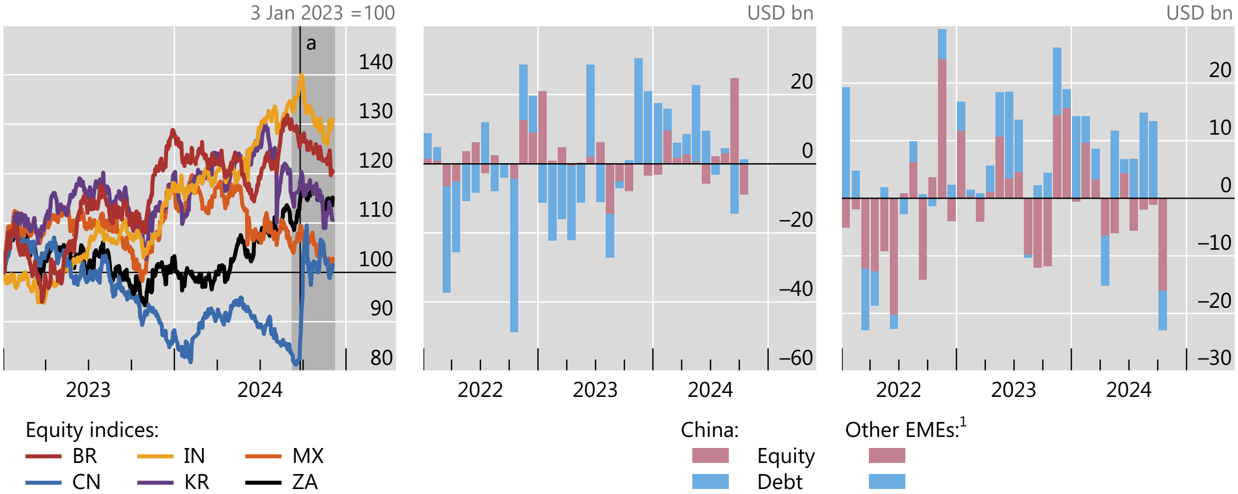


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[a] China stimulus announcement (24 September 2024).

[1] See technical annex for details.

Sources: Bloomberg; IIF; BIS.

EME portfolio flows at times diverged, while remaining volatile. In China, the surge in equity markets at the beginning of the review period went hand in hand with a sharp, albeit brief, rebound in equity inflows (Graph 12.B). Inflows into other EMEs concentrated mostly in portfolio debt early in the review period (Graph 12.C), but were also volatile, posting sharp outflows across debt and equities in October.

Stronger macroeconomic headwinds in China were also a key driver of commodity prices. The prices of agricultural commodities remained low, despite a brief surge at the time of the stimulus announcement. Oil prices were flat over the review period, notwithstanding a severe escalation of hostilities in the Middle East in October (Graph 13.A). Except for the price of iron ore, which did rise with the announcement, those of industrial metals also declined. Only gold and silver continued to climb, possibly reflecting geopolitical risks, but their prices then sank in the weeks following the US election (Graph 13.B).

On balance, the outcome in EMEs of rising yields, depreciating currencies and declining stock markets was a tightening of financial conditions. The tightening was overall much larger than in AEs and the United States in particular (Graph 13.C).

## Commodities reflect lower demand, financial conditions tighten

Graph 13

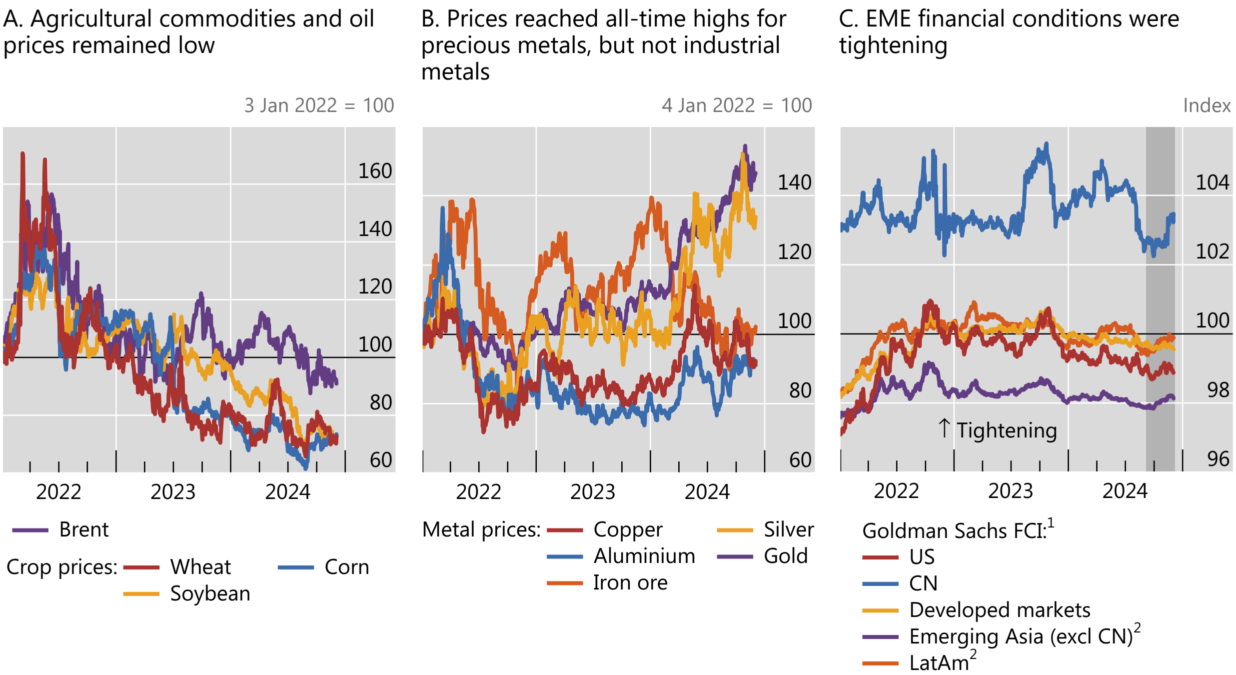


The shaded area indicates 7 September 2024–2 December 2024 (period under review).

[1] A value of 100 indicates average conditions. A higher (lower) value indicates tighter (looser) conditions. [2] See technical annex for details.

Sources: Bloomberg; Goldman Sachs Global Investment Research; national data; BIS.

## Negative interest rate swap spreads signal pressure in government debt absorption

*Matteo Aquilina, Andreas Schrimpf, Vladyslav Sushko and Dora Xia*①

Interest rate swaps are contracts in which counterparties agree to exchange a series of fixed rate payments for a series of floating payments linked to a benchmark rate. The swap *spread* is the difference between the interest rate swap rate and the yield on a government bond of the same maturity. The swap *rate* refers to the fixed rate in the swap; it reflects expectations of future floating rates and can therefore be interpreted as the price that ensures that both the receiver and the payer of the fixed rate view the contract as fairly priced from the start. Swap rates and bond yields are tied together by arbitrage. Absent costs involved in the arbitrage and compensation for risks, swap spreads should not deviate much from zero. Prior to the Great Financial Crisis, spreads were generally positive, as swap rates exceeded cash bond yields, in part due to some credit risk reflected in the swap rate.② More recently though, the constellation has flipped, with negative swap spreads being increasingly common across currencies and maturities.

### Interest rate swap spreads have shifted from positive to negative territory

In basis points | Graph B1

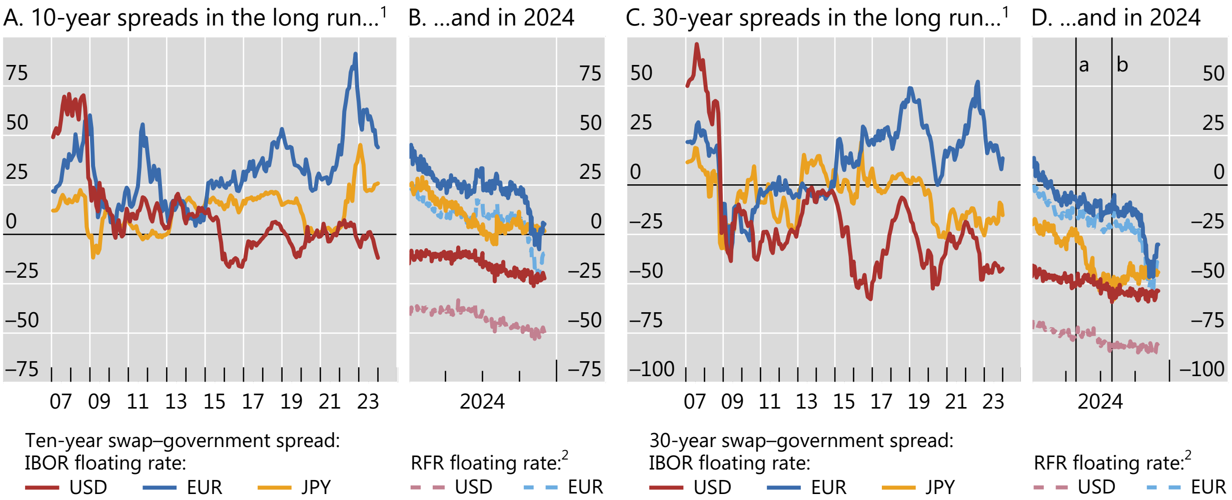


[a] Bank of Japan press conference (26 April 2024). [b] Bank of Japan press conference (31 July 2024).

IBOR = interbank offered rate; RFR= risk-free rate.

[1] Monthly averages of daily values. [2] Swap rates, and hence swap spreads, are lower when the floating rates are indexed to risk-free rates because the latter are overnight, so virtually risk-free. The effect is larger for US dollar swaps, because the RFR is a collateralised rate.

Sources: LSEG Workspace; authors' calculations.

Swap spreads in US dollars were negative for some time post-Great Financial Crisis (Graph B1)③ but have become more persistent even at 10-year tenors (Graph B1.B). More recently, negative swap spreads have also become more common in other major currencies, such as the euro and Japanese yen (Graph B1.D). Negative swap spreads, at least in theory, represent an arbitrage opportunity, and should therefore be quickly brought back to zero. Market participants could exploit a negative swap spread by holding a government bond funded through repo, paying the fixed rate in the swap and earning the floating, thus pocketing the difference between these rates.④

Negative swap spreads are not arbitraged away because they capture intermediation costs rather than a "free lunch". Negative spreads compensate intermediaries for holding government bonds on their balance sheets and entering swaps as fixed rate payers. Both swap and bond markets are intermediated by bank-affiliated dealers who require remuneration for using their balance sheets and taking on associated risks. When dealers absorb a large amount of bonds, they incur funding costs in the repo market for financing the long bond position. Additionally, they tend to hedge the interest rate risk by paying the fixed swap rate and receiving the floating rate. When doing so, dealers also need to factor in balance sheet costs from internal risk management and prudential rules, as well as opportunity costs of other uses of their balance sheet capacity. If these costs are high enough, dealers will recoup them through a negative swap spread. Moreover, if dealer balance sheets are constrained, non-bank players such as hedge funds may need to be incentivised to step in, deploying repo leverage to assume similar positions as dealers.

Supply-demand imbalances in government bond markets — Graph B2

A. Higher yields reflect "soft" auctions | B. Bonds on dealer balance sheets | C. Government bonds' funding costs

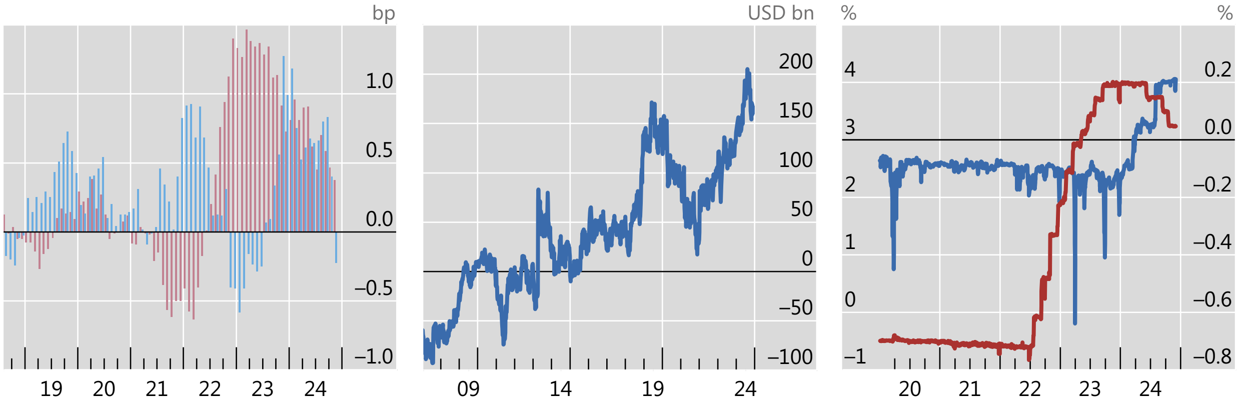


10-year Treasury auction yield surprise:[1] 10-year tail; 30-year tail

Dealers' net position in US Treasuries[2]

One-week repo rate: EUR (lhs); JPY (rhs)

[1] Auction yield minus when-issued yield; 12-month moving average. [2] Net US Treasury positions of primary dealers, across all tenors.

Sources: Bloomberg; LSEG Datastream; authors' calculations.

Downward pressure on swap spreads can originate either from a greater supply of government bonds in cash markets or a greater demand to receive the fixed rate in swap markets. In cash bond markets, investors' inability or unwillingness to absorb debt issuance or sales by other bond holders at prevailing prices exerts upward pressure on bond yields, pushing swap spreads lower. In swap markets, asset managers, institutional investors and corporate debt issuers are typical receivers of the fixed rate in the swap. The greater these players' demand to receive the fixed rate, the greater the downward pressure on swap rates and hence swap spreads.⑤

In recent months, pressure points primarily originating in the cash bond markets have caused noticeable imbalances in supply and demand. Recent Treasury auction results underscore this, with some monthly auctions clearing at a higher yield (lower price) than expected since 2022 (Graph B2.A). These "soft" auction results reflect investors' lukewarm interest in absorbing the supply of government debt and that dealer bond inventories have swollen to record levels this year (Graph B2.B).

Negative swap spreads have also emerged more recently in the euro area and Japan. A common driver with the United States is the additional expected bond issuance given the upward trajectory in debt supply. But jurisdiction-specific drivers are also at play. For one, as the euro area and Japanese central banks embarked on quantitative tightening (QT), expectations that the private sector would need to play a bigger marginal role in absorption pushed bond yields higher, compressing swap spreads. The 30-year yen swap spread dropped into deeply negative territory in April, when the Bank of Japan announced the exit from its ultra-loose balance sheet policy. The spread then decompressed in July, when the central bank revealed that the pace of QT would be slower than previously expected. In the euro area, swap spreads tumbled in October, as QT's effects were felt in markets via pressure on bond yields. And the steep fall in euro swap spreads coincided with a drop in the demand to tap the ECB Securities Lending Facility to source collateral (see main text) – a sign that government bond collateral in private hands was no longer scarce.

Another common driver of swap spreads has been the rise in funding costs in repo markets as central banks abandoned negative policy rates. In this case, the ECB was ahead of the Bank of Japan by almost two years (Graph B2.C). Since then, dealers have had to pay *positive* repo rates to fund their long positions in government bonds. Furthermore, the relative abundance of government bond collateral, due to increased issuance and less central bank heft, has put additional upward pressure on repo rates.

### The shift away from Libor mechanically lowered swap spreads

In basis points | Graph B3

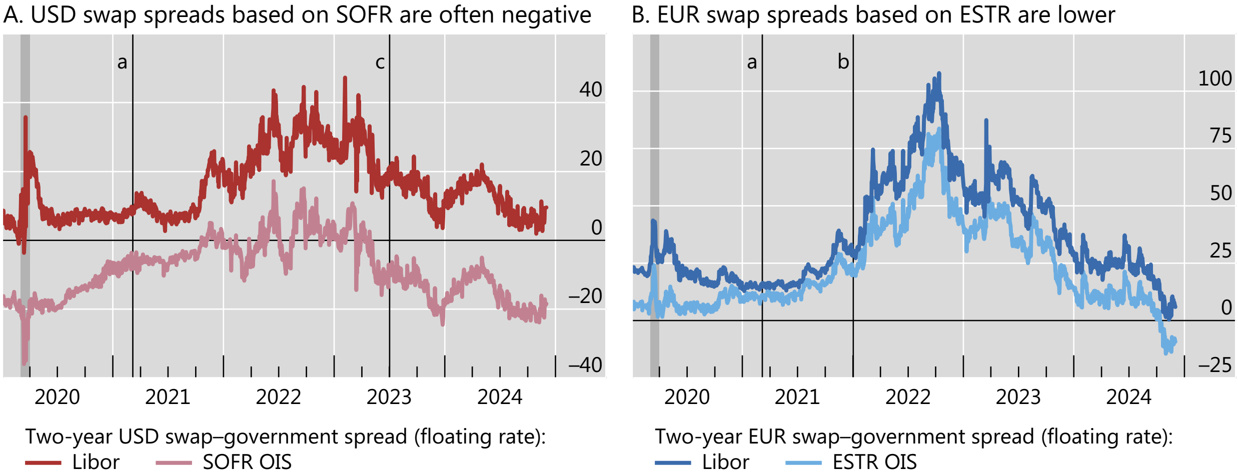


The shaded area indicates 1–31 March 2020 (Covid-19 financial turmoil).

ESTR = euro short-term rate; OIS = overnight index swap; SOFR = secured overnight financing rate.

[a] UK Financial Conduct Authority announces cessation of Libor (5 March 2021). [b] EUR Libor is no longer published (1 January 2022). [c] USD Libor is no longer published (1 July 2023).

Sources: LSEG Workspace; authors' calculations.

Finally, there are technical reasons for more negative swap spreads, notably the shift in reference rates in swaps away from Libor benchmark interest rates to so-called risk-free rates. Before the switch away from Libor, the swap rate was the fair price at inception of a series of expected future rates that embodied credit risk. Now that swap contracts are referenced to (nearly) risk-free rates in the floating leg, the swap rate itself is lower, which means that swap spreads also will be lower or more negative (Graphs B1.B and B1.D). In fact, the emergence of negative spreads at shorter tenors (eg two years) is mainly due to the shift towards risk-free benchmarks following the cessation of Libor (Graph B3). Benchmarks based on collateralised rates, such as the secured overnight financing rate (SOFR), result in even lower swap rates, and thus more negative swap spreads (Graph B3.A). SOFR-based spreads will also be influenced by repo market developments, as seen in the March 2020 dash for cash episode.⑦, ⑧

① The views expressed are those of the authors and do not necessarily reflect the views of the BIS. ② Credit risk played a bigger role when swap rates were linked to benchmarks based on unsecured term bank funding rates (eg Libor) and because counterparties in the swap were exposed to each other's credit risk (typically mitigated through collateralisation) – a risk significantly reduced through the shift to central clearing. ③ S Sundaresan and V Sushko, "Recent dislocations in fixed income derivatives markets", *BIS Quarterly Review*, December 2015. ④ See eg N Boyarchenko, P Gupta, N Steele and J Yen, "Negative swap spreads", Federal Reserve Bank of New York *Economic Policy Review*, October 2018; and U Jermann, "Negative swap spreads and limited arbitrage", *The Review of Financial Studies*, vol 33, no 1, 2019. ⑤ See S Klingler and S Sundaresan, "An explanation of negative swap spreads: demand for duration from underfunded pension plans", *The Journal of Finance*, vol 74, no 2, 2018; and, S Hanson, A Malkhozov and G Venter, "Demand-and-supply imbalance risk and long-term swap spreads", *Journal of Financial Economics*, vol 154, 103814, 2024. ⑥ The added sensitivity of long-maturity bonds to interest rate risk is another reason why negative swap spreads tend to be more prevalent at longer maturities. ⑦ See A Schrimpf and V Sushko, "Beyond LIBOR: a primer on the new benchmark rates", *BIS Quarterly Review*, March 2018; and D Wu and R Jarrow, "The Treasury – SOFR swap spread puzzle explained", 2024, available on SSRN. ⑧ Since the discontinuation of Libor, the difference between swap spreads based on the two types of rates is constant, as the ISDA fallback solution for legacy cleared derivatives transactions consists of a constant add-on risk spread.

# Technical annex

Graphs 8.A and 8.B: For Asia, ICE Asian Dollar indices.

Graph 8.C: US high-yield debt default rate based on Moody's 12-month rolling US speculative grade default rates. High-yield issuance based on a 12-month rolling sum.

Graph 9.A: Simple average of Bancorp, Citizens Financial Services, Dime Community Bancshares, Flagstar Financials, FVC Bankcorp, Live Oak Bancshares, Peapack-Gladstone Financial, Servisfirst Bancshares, Uscb Financial Holdings and Valley National.

Graphs 12.B and 12.C: Other EMEs = BR, CL, ID, IN, KR, LK, MN, MX, MY, PH, PK, TH, TW and VN.

Graph 13.C: For emerging Asia, GDP-PPP weighted average of ID, IN, KR, MY, PH and TH; for LatAm, GDP-PPP weighted average of BR, CL and MX.

Boris Hofmann
boris.hofmann@bis.org

Cristina Manea
cristina.manea@bis.org

Benoît Mojon
benoit.mojon@bis.org

# Targeted Taylor rules: monetary policy responses to demand- and supply-driven inflation[1]

*This feature documents that central banks operating under inflation targeting or similar regimes have in practice pursued their objectives in a targeted manner in the sense that they have reacted more forcefully to demand-driven than to supply-driven inflation. This new finding comes from the estimation of Taylor-type monetary policy rules for seven major advanced economies. The estimated targeted response aligns with both monetary theory prescriptions and central banks' doctrine as reflected in their official statements. Our analysis further suggests that during the post-pandemic inflation surge, policy rates were initially slow to respond but eventually caught up with the levels predicted by the targeted Taylor rules.*

*JEL classification: E12, E3, E52*

Mainstream monetary theory prescribes an asymmetric response to fluctuations in inflation, depending on their underlying drivers. If inflation stems from demand factors, the theory calls for a strong response to stabilise both inflation and output. By contrast, if inflation is due to supply factors, the central bank should give more weight to economic activity and partly "look through" associated inflationary pressures as long as inflation expectations remain firmly anchored (eg Erceg et al (2000); Bodenstein et al (2008)).

The doctrine of central banks operating under flexible inflation targeting or similar regimes reflects these principles. Central banks commonly acknowledge a more muted response to supply-driven inflation due to the often transitory nature of underlying shocks and to the induced macroeconomic trade-offs. As supply shocks push prices and output in opposite directions, central banks tend to respond more mildly to supply-driven inflationary (disinflationary) pressures to avoid reinforcing economic downturns (overheating). Central banks' common approach to accommodate supply-driven (dis)inflation underpins their medium-term price stability objectives.

[1] The views expressed do not necessarily reflect those of the Bank for International Settlements. We thank Ryan Banerjee, Claudio Borio, Matthieu Chavaz, Gaston Gelos, Marco Lombardi, Tsvetelina Nenova, Daniel Rees, Tom Rosewall, Damiano Sandri, Andreas Schrimpf, Hyun Song Shin and John Williams for helpful comments and suggestions. We are also grateful to Jose María Vidal Pastor for excellent research assistance.

*Key takeaways*

- *Over recent decades, central banks in advanced economies have responded much more strongly to demand- than to supply-driven inflation as gleaned from estimated targeted Taylor rules.*
- *This new finding corroborates both prescriptions of monetary theory and central bank doctrine as reflected in official statements.*
- *During the post-pandemic inflation surge, policy rates were initially slow to respond but eventually caught up with the levels predicted by the targeted Taylor rules.*

That said, conventional policy rules used to describe monetary policy reaction functions do not account for this asymmetric response. Instead, these rules essentially assume that policy rates respond to inflation in a uniform way irrespective of its underlying drivers (eg Taylor (1993); Clarida et al (2000); Smets and Wouters (2007)). In more refined versions of such rules, the lack of response to supply shocks due to their transitory nature is sometimes taken into account by formulating the policy rule in terms of core inflation (Carvalho et al (2021)) or an inflation forecast (Svensson (1997)).

This special feature bridges the gap between theory and doctrine on the one hand and the conventional description of policy reaction functions on the other.[2] We do so by estimating targeted Taylor rules that allow for different (targeted) responses to demand- and supply-driven inflation. To achieve this, we build on recent contributions by Shapiro (2022) and Eickmeier and Hofmann (2022) that disentangle the demand and supply factors underlying observable inflation dynamics. These approaches identify the time series of demand- and supply-driven inflation based on the basic conceptual consideration that a demand shock moves inflation and output

Decomposition of inflation in demand and supply factors[1]

Headline inflation, year on year, in per cent

Graph 1

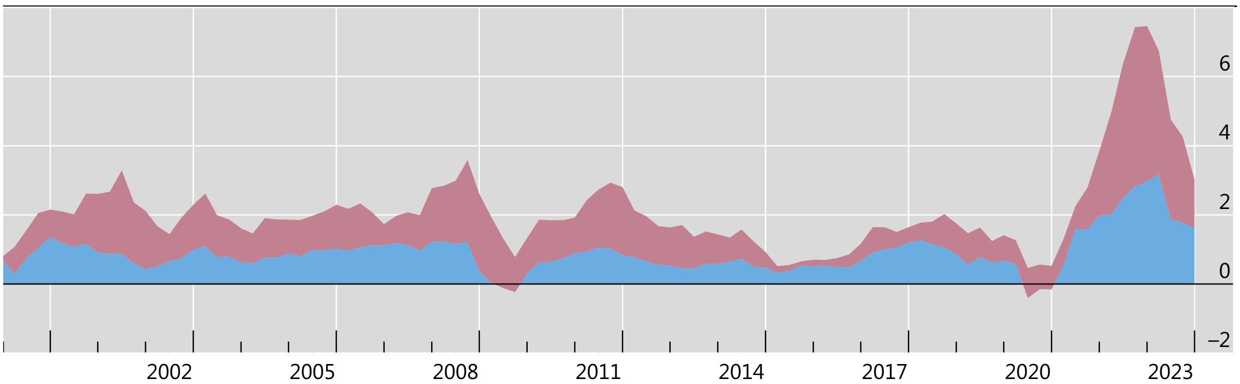


[1] Average across AU, CA, EA, GB, KR, SE and US. See Graph A.1 in the annex for the underlying individual series.

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); OECD; authors' calculations.

[2] The present analysis extends the empirical analysis in Hofmann et al (2024) by providing evidence beyond the United States. See the latter reference for a discussion of the implications of our results for the transmission of business cycle shocks and normative aspects pertaining to targeted versus conventional (unconditional) Taylor rules.

in the same direction, while a supply shock moves them in opposite directions. The resulting time series of demand- and supply-driven inflation highlight the changes in inflation drivers over time and, notably, reveal that the post-Covid-19 pandemic inflation surge was due to both demand and supply factors (Graph 1). We use these time series to estimate targeted Taylor rules for seven major jurisdictions (Australia, Canada, the euro area, Korea, Sweden, the United Kingdom and the United States) over periods when these economies operated price stability-oriented monetary policy regimes such as flexible inflation targeting.

Our core finding is that central banks have responded strongly to demand-driven inflation, and only weakly, if at all, to supply-driven inflation. For our baseline sample, the estimated response to demand-driven inflation is more than four times greater than that to supply-driven inflation. The difference is both statistically and economically significant. Thus, our findings suggest that central banks have pursued their price stability mandates in a *targeted* fashion, very much in line with both the prescriptions of monetary theory and central bank doctrine. Our results also imply that conventional monetary policy reaction functions are mis-specified, essentially estimating an average of the strong reaction of central banks to demand-driven inflation and their muted reaction to supply-driven inflation.

The feature proceeds as follows. The first section reviews the prescriptions of monetary theory and the central bank doctrine on how monetary policy should respond to demand- and supply-driven inflation. The second section describes the empirical analysis and compares our estimates of targeted Taylor rules to those of conventional Taylor rules. The third section sheds some new light on central bank responses during the post-pandemic inflation surge through the lens of our novel estimated targeted Taylor rules. The final section concludes and discusses potential challenges ahead for targeted policy reaction functions.

## Theory and central bank doctrine

Several theoretical considerations call for a more muted policy response to supply- than to demand-driven inflation. At the same time, official central bank communication suggests that these considerations also underpin central banks' doctrine with respect to interest rate setting.

### Theory

Two main reasons justify an asymmetric response to demand- versus supply-driven inflation.

First, certain types of supply shocks, like supply-driven commodity price shocks, tend to be transitory (Avalos et al (2025)). Despite their significant short-term impact on headline inflation, these shocks typically do not lead to persistently higher inflation (unless they induce large second-round effects on prices in other economic sectors). Because of long monetary policy transmission lags, responding to the uptick in inflation caused by such transient disturbances would be destabilising. By the time monetary policy would start to have traction, typically with a lag of 12 months or more, the effects of the shock on inflation would have largely subsided. Thus, reacting to the inflationary effects of such shocks would end up pushing inflation down too late and might risk unnecessarily weakening the economy (Mishkin (2007); Bandera et al (2023); Guerrieri et al (2023)).

A second reason for "looking through" supply-driven inflation arises from the specific stabilisation trade-off between inflation and real activity associated with supply shocks.[3] This can be illustrated with a simple aggregate demand (AD)–aggregate supply (AS) diagram (Graph 2). An inflationary demand shock (Graph 2.A) pushes both economic activity and prices up, ie in the same direction (shift of AD from (a) to (b)). Tighter monetary policy can thus counteract the shock's effects, essentially shifting the demand curve back (AD shift from (b) to (a)).[4] By contrast, a supply shock (Graph 2.B) gives rise to a trade-off between stabilising inflation or output. Specifically, an inflationary supply shock shifts the aggregate supply curve up and to the left, pushing prices up and output down (AS shift from (c) to (d)).[5] When monetary policy tightens to contain the rise in prices, aggregate demand shifts to the left from AD1 to AD2, reducing output even further and potentially leading to a severe recession (AD shift from (d) to (e)). For this reason, in standard macroeconomic models, the response of monetary policy to inflationary supply shocks is to tighten, but not as much as for demand shocks, especially if there is a chance that the shock itself (or its impact on inflation) is temporary (eg Erceg et al (2000); Blanchard and Galí (2007); Bodenstein et al (2008)).[6]

Monetary policy trade-offs in the face of supply shocks | Graph 2

A. Monetary policy can counteract the effects of inflationary demand shocks on both inflation and output

B. Monetary policy response to inflationary supply shocks further dampens economic activity

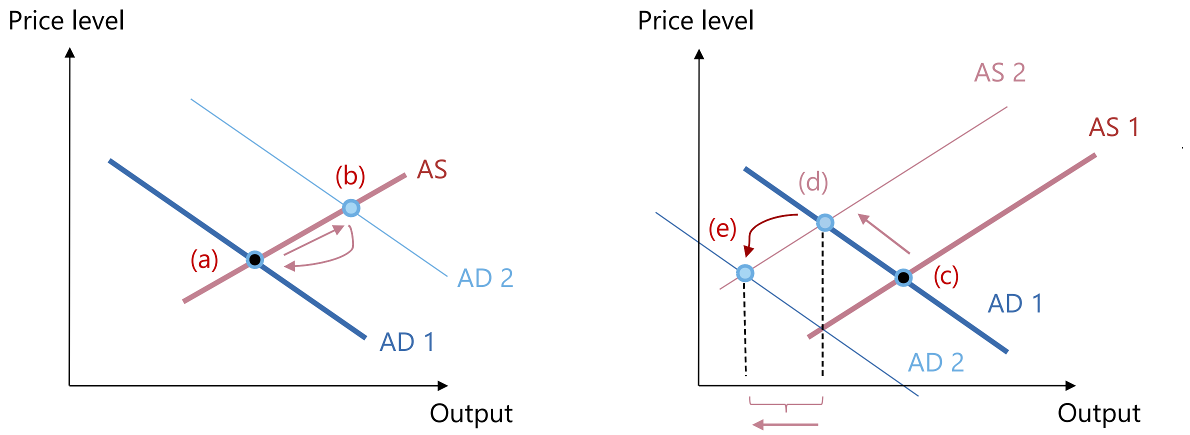


AD = aggregate demand; AS = aggregate supply.

Source: Authors' elaboration.

[3] This is the macroeconomic stabilisation trade-off applying to business cycle frequency. Central banks may face an additional trade-off applying to lower time frequencies between price stability and financial stability. It may arise in particular from very persistent favourable supply shocks and may constitute a third reason to look through supply-driven inflation (see Borio (2006) for the case of the disinflationary effects of globalisation).

[4] The lack of a trade-off between stabilising inflation and stabilising output around its desired level is known in the literature as the "divine coincidence" (Blanchard and Galí (2007)). According to monetary theory, this property generally holds for demand shocks, but not for supply shocks (Galí (2015)).

[5] Typical inflationary supply shocks are increases in markups or in the price of energy, or decreases in the pace of productivity.

[6] The same logic applies to disinflationary supply shocks (eg an acceleration of productivity or a decrease in the price of energy), which push output up and inflation down potentially below target.

A necessary condition for central banks' ability to look through is that longer-term inflation expectations are firmly anchored to central banks' inflation targets. Well anchored inflation expectations limit the risk that the initial impact of supply shocks is amplified through wider price and wage adjustments, preventing supply-driven inflation from turning into wage-price spirals. By contrast, when there is a risk that inflation expectations de-anchor, central banks are called upon to react forcefully to inflation irrespective of its underlying drivers (eg Reis (2022); Bandera et al (2023)).[7]

Central bank doctrine: official communications

Table 1

| Institution | Communications |
|---|---|
| Reserve Bank of Australia | • A central bank may "look through" the price effects of a supply shock if it is expected to be short-lived and inflation expectations remain anchored (RBA (2023)).<br>• Life is more complicated in a world of supply shocks; an adverse supply shock increases inflation and reduces output and employment (Lowe (2022)).<br>• If inflation expectations do increase and wage- and price-setting behaviour responds to the higher inflation, an interest rate response is required (Lowe (2023)). |
| Bank of England | • The orthodox monetary response to a global shock to energy prices is to "look through" them.<br>• When the economy is hit by temporary cost shocks, policymakers face a trade-off [between] output and inflation (Tenreyro (2022)).<br>• If inflation expectations drift away, monetary policy needs to lean against inertia to return inflation to target (Bandera et al (2023)). |
| Federal Reserve Board | • Standard monetary prescription is to "look through" commodities price shocks (Brainard (2022)).<br>• The response to the inflationary effects of supply shocks should be attenuated. Supply shocks tend to move prices and employment in opposite directions (Powell (2023)).<br>• Supply shocks that drive inflation high enough can affect the longer-term inflation expectations. Monetary policy must forthrightly address risks of de-anchoring of expectations (Powell (2023)). |
| European Central Bank | • When faced with supply shocks, central banks can, in principle, "look through" them, as these shocks will usually leave no lasting imprint on inflation.<br>• The appropriate policy response will depend on the type of shock. For a supply shock, price stability may conflict with the contractionary impact of the shock. (Papademos (2003)).<br>• In situations where inflation expectations can de-anchor, central banks must then react forcefully to prevent above-target inflation becoming entrenched (Lagarde (2024)). |
| Bank of Canada | • The bank's framework for inflation targeting allows temporary supply shocks to be largely ignored, as long as they do not feed into inflation expectations (Dodge (2002)).<br>• Supply shocks present central banks with a difficult trade-off between growth and inflation. We focus on balancing the upside risks to inflation with the downside risks to growth (Macklem (2024)). |
| Sveriges Riksbank | • Supply shocks such as shortage of snow that restricted the supply of hydroelectric power can occasion deviations from the inflation target (Bäckström (2002)).<br>• Supply shocks present a challenge: policymakers want to prevent inflation from becoming entrenched at a high level but want to avoid exacerbating the downturn. (Thedéen (2023)).<br>• If there is a risk that inflation exceeds 2 percent for a long time, a tighter monetary policy may be necessary to maintain confidence in the inflation target (Löf and Stockhammar (2024)). |
| Bank of Korea | • If inflation is projected to exceed the target but the real sector faces supply shocks, the central bank should decide whether to adjust interest rates to ensure price stability (Bank of Korea (2017)). |

Sources: Central bank statements; authors' elaboration.

In that case, monetary policy would loosen but not as much as in case of a disinflationary demand shock, as in the case of the latter output is below its potential.

7 See also Maechler (2024) for a recent discussion of this point.

### Central bank doctrine

Central bank doctrine – according to central bank official statements – aligns with these conceptual prescriptions (Table 1). All of the seven central banks included in our analysis refer to the transitory nature of many supply shocks and to the stabilisation trade-offs between inflation and real activity associated with such shocks as reasons to at least partially look through them. They also typically refer to the anchoring of inflation expectations as a precondition for their ability to do so.

## Empirical analysis

Are the prescriptions of monetary theory and the stated central bank doctrine also reflected in practice, ie in the actual conduct of monetary policy? If this were the case, the monetary policy response to demand-driven inflation would be measurably stronger than that to supply-driven inflation. This is an empirical question, which we address by estimating monetary policy reaction functions.

Our analysis draws on recent empirical contributions providing a decomposition of inflation into demand- and supply-driven components. Specifically, we use the decompositions of inflation based on the methods proposed by Shapiro (2022) for Australia, Canada, Korea, Sweden, the United Kingdom and the United States, and by Eickmeier and Hofmann (2022) for the euro area; see Box A for details.[8]

We determine the sample period according to the availability of the inflation decomposition and the beginning of price stability-focused monetary policy regimes in each of the seven jurisdictions covered by our analysis. Our sample includes observations starting in the third quarter of 1979 (when Paul Volcker was appointed Chair of the Federal Reserve) for the United States, post-2002 for the euro area, post-2000 for the United Kingdom and post-1999 for the other jurisdictions. The sample ends with the most recently available observation, respectively, ranging between the second quarter of 2024 for the US and the first quarter of 2023 for Korea (see annex B for details). To address the challenge posed by the (proximity to the) zero lower bound, we include in our baseline specification only observations for which the policy rate was above 0.5%.[9]

8 The analysis covers the seven jurisdictions for which such decompositions are available. In the case of the United States, for which both decompositions are available, we use the series derived with the method in Shapiro (2022) as our baseline, and then check that our findings carry through when using the series based on the method in Eickmeier and Hofmann (2022).

9 Our findings also hold if we use Wu and Xia (2016) shadow rates when available (for the US, the euro area and the UK). Under this specification, the estimated coefficient of demand-driven inflation is slightly higher compared with our baseline, while that of supply-driven inflation does not materially change. Moreover, we assessed the sensitivity of the results to the calibration of the threshold for the lower bound. When choosing a lower bound below 0.5% under our baseline specification, the estimated monetary policy responses to demand-driven and supply-driven inflation slightly decrease, consistent with the limited policy space characterising the additional observations included in the estimations in these cases compared with our baseline. These results are available upon request.

Box A

## Decomposing inflation into demand and supply components

To decompose inflation into its demand- and supply-driven components, we primarily rely on the approach of Shapiro (2022), complemented with that of Eickmeier and Hofmann (2022). Both methods are based on sign restrictions motivated by a standard supply and demand framework. That is, changes in supply move inflation and output in opposite directions, while changes in demand move both variables in the same direction.

Shapiro (2022) decomposes personal consumption expenditure (PCE) inflation in two time series, the supply- and demand-driven contributions, which quantify the degree to which either demand or supply is driving inflation in each month. The identification relies on the signs of the residuals from separate price and quantity regressions for each of the more than 100 goods and services categories in the PCE price index. Specifically, for each category current price (quantity) levels are regressed on their own lagged values and lagged values of quantity (price) levels. Categories for which the residuals of the two regressions have the *same sign* in a given month are classified as subject to a demand shock in that month; the corresponding price change is then labelled as *demand-driven*. By contrast, categories with residuals of *opposite signs* in a given month are classified as subject to a supply shock in that month; the corresponding price change is then labelled as *supply-driven*. The demand-driven (supply-driven) contribution to inflation in a given month is then measured as the expenditure-weighted average of the price changes of those categories labelled as demand-driven (supply-driven) in that month. Decompositions from the Shapiro approach are available for the Australia, Canada, Korea, Sweden, the United Kingdom and United States, but not for the euro area.

### Demand/supply inflation decomposition for the US

Headline inflation, year on year, in per cent

Graph A1

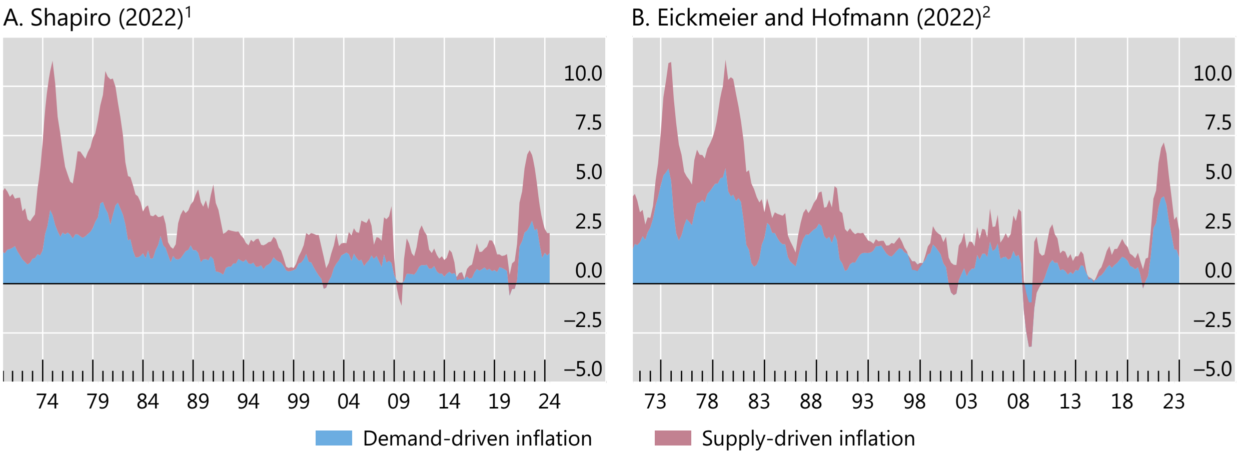


[1] Quarterly average of monthly series. [2] The original Eickmeier and Hofmann (2022) inflation decomposition is expressed in quarter-on-quarter changes, and refers to standardised inflation series. For our estimations, we use the mean and standard deviation of the aggregate inflation series to back out the demand and supply components of inflation, and then express those two components as year-on-year changes.

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); authors' calculations.

We complement the data from the Shapiro approach with the demand-supply inflation decomposition for the euro area based on Eickmeier and Hofmann (2022). This approach estimates indicators of demand- and supply-driven inflation based on a factor model using more than 140 quarterly time series of inflation and real activity measures. The estimation relies on imposing sign restrictions on factor loadings. Supply is identified as a factor that loads negatively on inflation and positively on economic activity; demand, in turn, is identified as a factor that loads positively on both inflation and economic activity.

While the two approaches are different, they yield broadly similar decompositions of inflation into demand and supply components. This is what a comparison suggests for the United States, the only economy for which both decompositions are available (Graph A.1). For instance, for the post-pandemic inflation surge, the two decompositions indicate that both demand and supply forces contributed to the inflation surge, with the Shapiro approach attributing a somewhat larger role to supply factors. More generally, the correlation coefficient of the demand (supply) series derived with the two methods is in the ballpark of 86% and is highly statistically significant.

We consider two empirical specifications of the Taylor rule: a conventional one (our benchmark) whereby the central bank is assumed to care about fluctuations in inflation and real activity, and a targeted one, which allows for a different response to the demand- and supply-driven components of inflation.

The benchmark specification for the monetary policy reaction function – as described by a conventional Taylor rule with interest rate smoothing, takes the following form:

$$i_t = i + \rho i_{t-1} + (1-\rho)[\alpha\pi_t + \beta\, y_t] + \varepsilon_t \qquad (1)$$

where $i_t$ is the (annualised) policy rate in quarter $t$, $\pi_t$ is (year-on-year) inflation, $y_t$ is the output gap[10] and $i$ is a constant reflecting the long-term level of interest rates.[11] The lagged interest rate $i_{t-1}$ captures central banks' proclivity for smoothing adjustments in the policy rate over time.

The targeted Taylor rule is an extension of the benchmark specification in (1), in which we replace overall inflation $\pi_t$ with its demand- and supply-driven components $\pi_t^d$ and $\pi_t^s$ , where $\pi_t = \pi_t^d + \pi_t^s$:

$$i_t = i + \rho i_{t-1} + (1-\rho)[\alpha^d\pi_t^d + \alpha^s\pi_t^s + \beta\, y_t] + \varepsilon_t \qquad (2)$$

## Conventional and targeted Taylor rules: panel estimates[1]

In percentage points — Graph 3

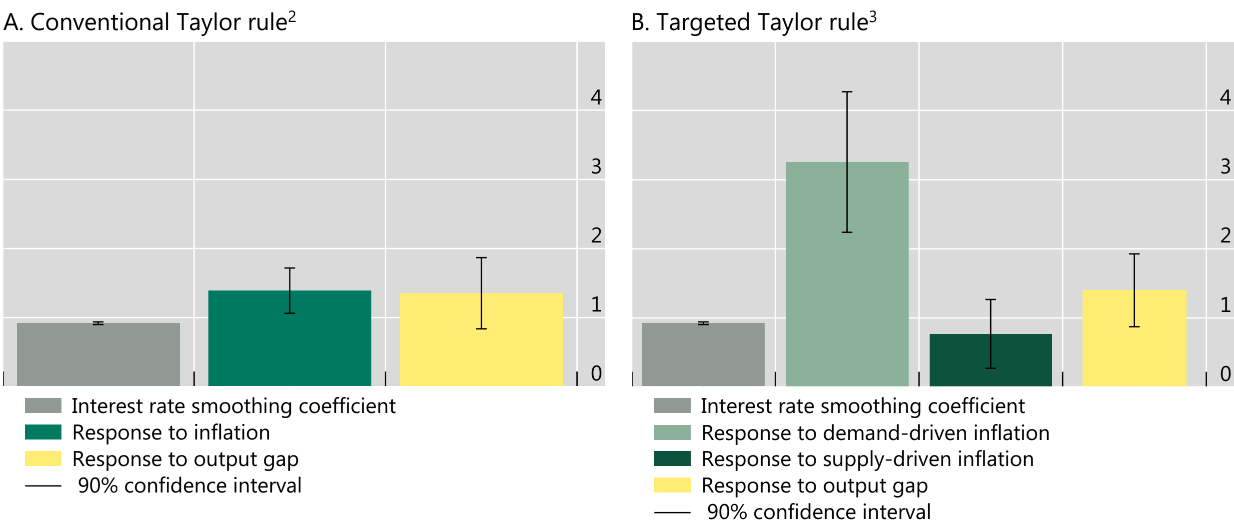


[1] Panel estimates for AU, CA, EA, GB, KR, SE and US. [2] Estimated coefficients of $\rho$, $\alpha$, $\beta$ in the conventional Taylor rule in equation (1). [3] Estimated coefficients of $\rho, \alpha^d, \alpha^s, \beta$ in the targeted Taylor rule in equation (2).

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); OECD; national data; authors' calculations.

[10] As measures of the output (real GDP) gap, we use national series from central bank sources whenever available (ie for the United States, Canada and Australia) and standard two-sided Hodrick-Prescott (HP) filtered measures of the output gap otherwise. Our findings carry over when using the two-sided HP filtered measures for all jurisdictions. As we use the latest vintage of data available – and not the data available in real time to policymakers – our estimates provide de facto an ex post description of central banks' monetary policy reaction functions.

[11] In the simplified way the Talyor rule is specified in equation (1), the constant $i$ is an amalgam of the long-term level of the real interest rate, the central bank's inflation target and the model parameters.

The first set of estimates, reported in Graph 3, describes the average monetary policy reaction function of the central banks in focus.[12]

The results strongly support the targeted specification of the Taylor rule. For the standard Taylor rule from equation (1), the estimated inflation coefficient (α) equals 1.39, while that of the output gap (β) equals 1.35 (Graph 3.A). The coefficient of inflation is slightly below that prescribed by the original Taylor (1993) rule (1.38 instead of 1.5), while that of the output gap is higher (1.35 instead of 0.5) consistent with recent estimates for the post-Volcker period.[13] However, the estimates of the targeted specification of the Taylor rule from equation (2) suggest that there is a significant difference in the response to demand- versus supply-driven inflation (Graph 3.B). Notably, central banks appear to respond more forcefully to demand- than to supply-driven inflation (more than fourfold, 3.26 versus 0.77). The difference is both statistically and economically significant.

Individual estimates at the jurisdiction level of the conventional Taylor rule (Graph 4.A) and of the targeted Taylor rule (Graph 4.B) confirm this result. In particular, for the targeted Taylor rule the interquartile range of individual responses to demand-driven inflation ranges from 1.5 to 4.1 and lies on top of that of supply-driven inflation which ranges from –0.47 to 1.03.

How do our targeted Taylor rules compare with alternative ways of capturing the dependence of the monetary policy response on the underlying drivers of inflation? A common approach to model such a dependence is to specify the monetary policy reaction function in terms of core and non-core inflation, with the former capturing

Conventional and targeted Taylor rules: individual estimates across economies

In percentage points | Graph 4

A. Conventional Taylor rule[1]

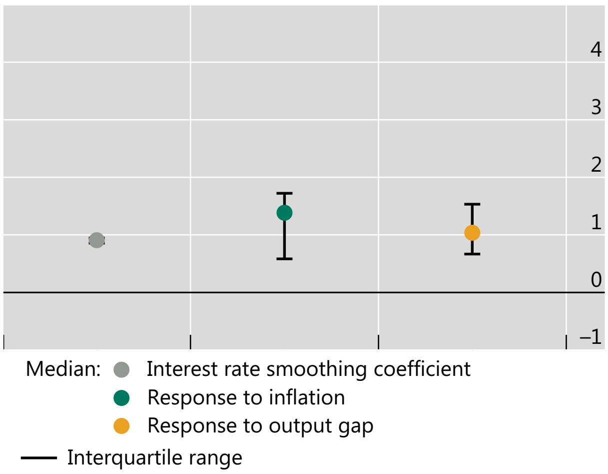


B. Targeted Taylor rule[2]

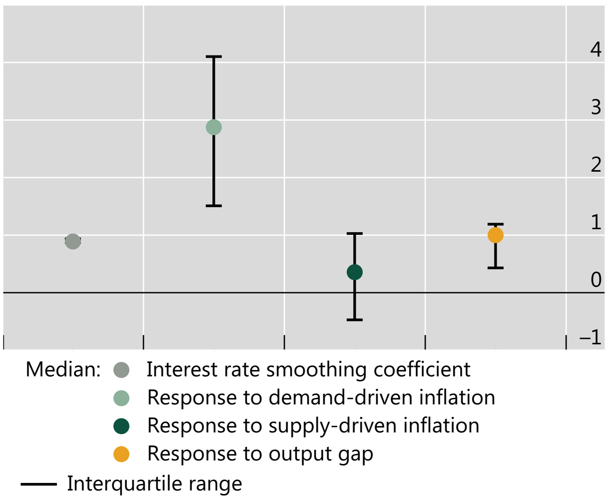


[1] Individual coefficient estimates of ρ, α, β in the conventional Taylor rule (1) estimated at jurisdiction level. [2] Individual coefficient estimates of $\rho, \alpha^d, \alpha^s, \beta$ in the targeted Taylor rule (2) estimated at jurisdiction level.

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); OECD; national data; authors' calculations.

[12] See annex B for further technical details on the estimation procedure.

[13] For example, Carvalho et al (2021) report a value of 0.81 on the output gap coefficient for the post-Volcker pre-Great Financial Crisis period for the United States.

developments in underlying inflation and the latter more transitory inflation fluctuations. There are different approaches to estimating core inflation (and non-core as the residual to headline inflation). One simple approach is to measure non-core inflation as the rate of change in the energy and food components of the consumer price index, and core inflation as the rate of change of the remainder of the index. To the extent that food and energy price changes are more transitory and often supply-driven, while underlying core inflation may often reflect primarily demand conditions, the core versus non-core inflation distinction may capture similar aspects as our demand- versus supply-driven inflation distinction.

We again estimate two different specifications of the monetary policy reaction function, focusing on the reaction of the US Federal Reserve to core inflation and its demand- and supply-driven components. We conduct this analysis only for the United States because the demand versus supply decomposition of core inflation is not available in other jurisdictions. The first specification splits inflation $\pi_t$ into its core inflation $\pi_t^c$ and non-core inflation $\pi_t^{nc}$, ie changes in food and energy prices:

$$i_t = i + \rho i_{t-1} + (1-\rho)[\alpha^c \pi_t^c + \alpha^{nc} \pi_t^{nc} + \beta\, y_t] + \varepsilon_t \qquad (3)$$

The second specification splits inflation into its core- and non-core components, and in addition the former into its demand- and supply-driven components:

$$i_t = i + \rho i_{t-1} + (1-\rho)\left[\alpha^{c,d} \pi_t^{c,d} + \alpha^{c,s} \pi_t^{c,s} + \alpha^{nc}\, \pi_t^{nc} + \beta\, y_t\right] + \varepsilon_t \qquad (4)$$

The coefficients $\alpha^{c,d}$ and $\alpha^{c,s}$ are now the coefficients on the demand and supply components of core inflation.

The results reported in Graph 5 suggest that the distinction between core and non-core inflation does not capture the distinction between demand- and supply-driven inflation. We find that the Federal Reserve responded strongly to core inflation, with a response coefficient of above 2, while it essentially ignored the energy and

## Taylor rules for the US: core versus non-core inflation

In percentage points — Graph 5

A. Taylor rule with core and non-core inflation[1] — B. Targeted Taylor rule with core and non-core inflation[2]

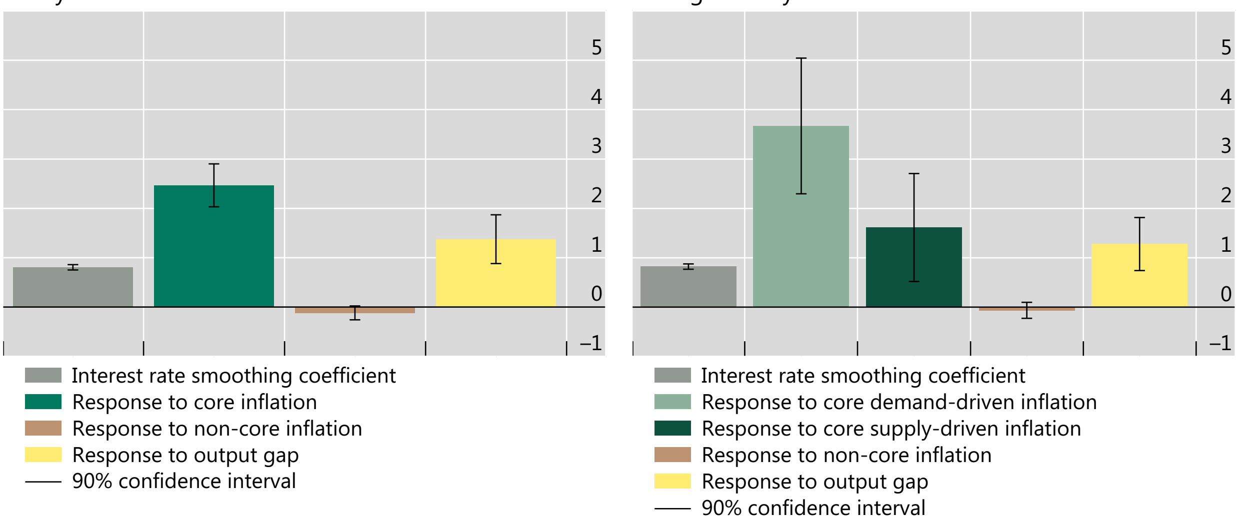


[1] Coefficient estimates of $\rho$, $\alpha^c$, $\alpha^{nc}$, $\beta$ in a conventional Taylor rule specification with core and non-core inflation. [2] Coefficient estimates of $\rho$, $\alpha^{c,d}$, $\alpha^{c,s}$, $\alpha^{nc}$, $\beta$ in a targeted Taylor rule specification with core and non-core inflation.

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); national data; authors' calculations.

food price changes (Graph 5.A), consistent with the pattern found for a broader group of advanced economies in Avalos et al (2025). However, the estimates of equation (4) suggest that the response to demand-driven fluctuations of core inflation was about three times stronger than that to supply-driven fluctuations (Graph 5.B). Thus, our result that the central bank reacted more strongly to demand-driven fluctuations in inflation holds also when focusing on core inflation.[14]

## Targeted Taylor rules and the post-pandemic inflation surge

What does our analysis have to say about the monetary policy reaction to the post-pandemic inflation surge? To answer this question, we compare the observed path of policy rates with those implied by the estimated targeted Taylor rules since the beginning of the inflation surge. The latter corresponds, for each jurisdiction, to the first quarter after Q3 2020 when inflation exceeded 2%.

The results of this exercise suggest that policy rates were initially slow to respond to inflation but then rapidly caught up with the levels implied by the estimated Taylor rules. Central banks set policy rates at lower levels than those predicted by the estimated rule in the first four quarters of the inflation surge (Graph 6.A, blue bars). By the fourth quarter, the policy rate was more than 50 basis points below the implied rate in half of the jurisdictions in focus. Subsequently, as both supply- and demand driven inflation took hold (red and blue line, respectively), central banks tightened more forcefully, bringing policy rates up to, and eventually even above, levels consistent with the targeted Taylor rules.

The initial slow reaction may at least partly reflect an initial misdiagnosis that the increase in inflation was primarily supply-driven at that stage, so that central banks could look through it.[15] Indeed, in late 2020 and throughout 2021 policy rates were close to the level implied by the estimated targeted rule if inflation were exclusively supply-driven (Graph 6.A, red line). However, there was a significant demand-driven component, reflected in the considerably higher level of policy rates predicted by the estimated targeted Taylor rules based on the actual decomposition of inflation (Graph 6.B, yellow line).

[14] We also checked whether the strong estimated response to demand-driven inflation is not equivalent to reacting to the inflation forecast. We add the consensus inflation forecast as an additional regressor in equation (2) and obtain highly statistically significant coefficients for all terms in our regressions, with the coefficient of supply-driven inflation turning negative. These results suggest that central banks account for both contemporaneous inflation and its forecast, and that everything else equal, react less when contemporaneous supply-driven inflation is high. This is consistent with central banks' concern of stabilisation trade-off between inflation and real activity in those latter cases. The results are available upon request.

[15] Other complementary explanations could also be relevant. One such example is that central banks delivered on the promises made under forward guidance (or flexible average inflation targeting in the United States) to overshot inflation targets once exiting the zero lower bound.

Post-pandemic inflation surge: deviations from the estimated targeted Taylor rule

In percentage points Graph 6

A. Deviations from estimated policy rule and inflation[1]

B. Policy rate paths: observed versus counterfactual[2]

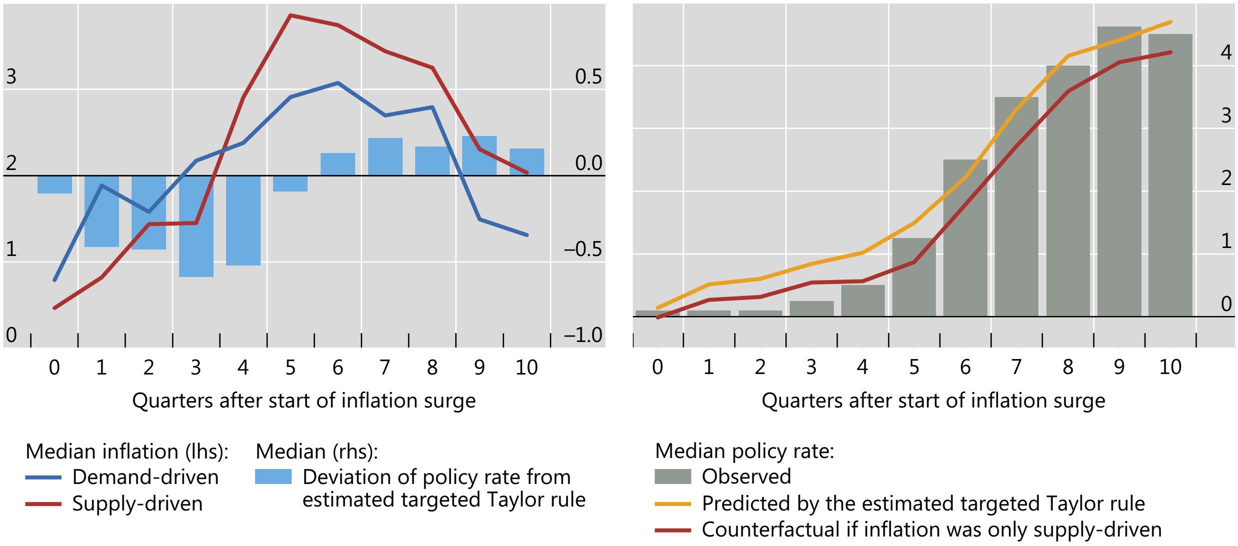


[1] The blue bars show the median deviation of the observed policy rate from the targeted Taylor rule estimates during the inflation surge. The beginning of the inflation surge in each jurisdiction is defined as the first quarter post-Q3 2020 when inflation exceeded 2%. [2] The yellow line shows the median predicted targeted policy rate path given the observed demand- and supply-driven inflation series. The red line shows the counterfactual policy rate path if the observed aggregate level of inflation was entirely supply-driven. Note that the difference between the median of the predicted policy rate (yellow line) and the median of observed policy rate (grey bars) is not equal to the median of the deviations of the observed policy rate from the estimated rule (panel A, blue bars).

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); OECD; national data; authors' calculations.

## Conclusions

The analysis in this special feature suggests that central banks operating under inflation targeting or similar regimes have conducted policy in a "targeted" manner. They did so by responding strongly to demand-driven inflation, but only mildly to supply-driven inflation. These findings are consistent with both central bank doctrine and the prescriptions of the prevailing monetary theory paradigm. They also suggest that conventional monetary policy reaction functions commonly used in academic research and central bank analytical models which assume a one-size-fits-all policy with respect to inflation are mis-specified.

An important question concerns the feasibility of targeted Taylor rules as a guide to monetary policy in real time. The measures of demand- and supply-driven inflation we have used became available only recently. Our findings therefore imply that central banks were able to infer similar information about the supply- versus demand-driven nature of inflation from their indicators, analytical toolboxes and judgment. However, the finding of a potential misdiagnosis of the nature of inflation at the beginning of the inflation surge suggests that this inference can be difficult in real-time conditions. The availability of the new indicators of demand- and supply-driven inflation could help mitigate this problem as they seem to have provided quite reliable decompositions in real time.[16]

Going forward, greater supply headwinds may force central banks to react more forcefully to supply disturbances to ensure inflation expectations stay firmly anchored. This may, in turn, constrain central banks' ability to conduct their policy in a targeted fashion. As recently highlighted by the central bank community (eg Bandera et al (2023); Maechler (2024)), geopolitical tensions and climate change could lead to larger, more frequent and more persistent adverse supply disturbances. At the same time, ageing populations and deglobalisation may render the supply side of the economy less elastic, magnifying the effect of business cycle shocks on inflation, irrespectively of their nature. If such changes materialise, they will call for a stronger reaction to supply-driven inflation to curb the second-round effects of supply shocks and prevent a de-anchoring of inflation expectations.

16 See the real-time analysis in Eickmeier and Hofmann (2022).

## References


Avalos, F, R Banerjee, M Burgert, B Hofmann, C Manea and M Rottner (2025): "Commodity prices and monetary policy: old and new challenges", *BIS Bulletin*, forthcoming.

Bäckström, U (2002): "Sweden's economy and monetary policy", speech at Skanes Provinsbank, Helsingborg, 7 March.

Bank of Korea (2017): *Monetary policy in Korea,* Bank of Korea, December.

Bandera, N, L Barnes, M Chavaz, S Tenreyro and L von dem Berge (2023): "Macroeconomic stabilisation in a volatile inflation environment", paper presented at the 2023 ECB Forum on Central Banking.

Blanchard, O and J Galí (2007): "Real wage rigidities and the New Keynesian model", *Journal of Money, Credit and Banking*, vol 39, pp 35–65.

Bodenstein, M, C Erceg and L Guerrieri (2008): "Optimal monetary policy with distinct core and headline inflation rates", *Journal of Monetary Economics*, vol 55, pp S18–S33.

Borio, C (2006): "Monetary and prudential policies at a crossroads? New challenges in the new century", *BIS Working Papers*, no 216, September.

Brainard, L (2022): "What can we learn from the pandemic and the war about supply shocks, inflation, and monetary policy?", speech at the 21st BIS Annual Conference "Central banking after the pandemic: challenges ahead", Basel, 28 November.

Carvalho, C, F Nechio and T Tristao (2021): "Taylor rule estimation by OLS", *Journal of Monetary Economics*, vol 124, pp 140–54.

Clarida, R, J Galí and M Gertler (2000): "Monetary policy rules and macroeconomic stability: evidence and some theory", *The Quarterly Journal of Economics*, vol 115, no 1, pp147–80.

Dodge, D (2002): "Macroeconomic stabilization policy in Canada", speech at the Federal Reserve Bank of Kansas City Jackson Hole symposium, 31 August 31.

Eickmeier, S and B Hofmann (2022): "What drives inflation? Disentangling demand and supply factors", *BIS Working Papers*, no 1047, forthcoming in *International Journal of Central Banking*.

Erceg, C, D Henderson and A Levin (2000): "Optimal monetary policy with staggered wage and price contracts", *Journal of Monetary Economics*, vol 46, no 2, pp 281–313.

Galí, J (2015): *Monetary policy, inflation, and the business cycle: an introduction to the New Keynesian framework* (2nd ed), Princeton University Press.

Guerrieri, V, M Marcussen, L Reichlin and S Tenreyro (2023): "The art and science of patience: relative prices and inflation", *Geneva Reports on the World Economy*, vol 26.

Hofmann, B and B Bogdanova (2012): "Taylor rules and monetary policy: a global 'Great Deviation'?", *BIS Quarterly Review*, September, pp 37–49.

Hofmann, B, C Manea and B Mojon (2024): "Targeted Taylor rules: some evidence and theory", *BIS Working Papers*, no 1234, December.

Lagarde, C (2024): "Monetary policy in an unusual cycle: the risks, the path and the costs", introductory speech at the opening reception of the ECB Forum on Central Banking, Sintra, 1 July.

Löf, M and P Stockhammar (2024): "What factors drove the surge in inflation?", Sveriges Riksbank *Staff Memos*, June.

Lowe, P (2022): "Price stability, the supply side and prosperity", speech at the CEDA Annual Dinner Address, Melbourne, 22 November.

——— (2023): "Monetary policy, demand and supply", speech at the National Press Club, Sydney, 5 April.

Macklem, T (2024): "Rewired, recast and redirected: global trade and implications for Canada", speech at the Canada-UK Chamber of Commerce, London, 10 September.

Maechler, A (2024): "Monetary policy in an era of supply headwinds: do the old principles still stand?", speech at the London School of Economics, London, 2 October.

Mishkin, F (2007): "Headline versus core inflation in the conduct of monetary policy", remarks at the "Business cycles, international transmission and macroeconomic policies" conference, HEC Montréal, 20 October.

Papademos, L (2003): "Economic cycles and monetary policy", speech at the. International Symposium of the Banque de France on "Monetary policy, the economic cycle and financial dynamics", Paris, 7 March.

Powell, J (2023): "Monetary policy challenges in a global economy", opening remarks at a policy panel at the 24th Jacques Polak Annual Research Conference, Washington DC, 9 November.

Reis, R (2022): "The burst of high inflation in 2021–22: how and why did we get here?", *BIS Working Papers*, no 1060, December.

Reserve Bank of Australia (RBA) (2023): "Statement on monetary policy – Box C: Supply and demand drivers of inflation in Australia", February.

Shapiro, A (2022): "Decomposing supply and demand driven inflation", Federal Reserve Bank of San Francisco *Working Papers*, no 2022-18, forthcoming in *Journal of Money, Credit and Banking*.

Smets, F and R Wouters (2007): "Shocks and frictions in US business cycles: a Bayesian DSGE approach", *American Economic Review*, vol 97, no 3, pp 586–606.

Svensson, L (1997): "Inflation forecast targeting: implementing and monitoring inflation targets", *European Economic Review*, vol 41, no 6, pp 1111–46.

——— (1999): "Inflation targeting as a monetary policy rule", *Journal of Monetary Economics*, vol 43, no 3, pp 607–54.

Taylor, J (1993): "Discretion versus policy rules in practice", *Carnegie-Rochester Conference Series on Public Policy*, vol 39, pp 195–214.

Tenreyro, S (2022): "The economy and policy trade-offs", 2022 Dow Lecture, National Institute of Economic and Social Research, London, 23 February.

Thedéen, E (2023): "Lessons from a turbulent period", speech at the Carnegie Fonder, Stockholm, 20 December.

Wu, J C and D Xia (2016): "Measuring the macroeconomic impact of monetary policy at the zero lower bound", *Journal of Money, Credit, and Banking*, vol 48, no 2–3, pp 253–91.

# Annex A

## Decomposition of inflation in demand and supply factors: baseline specification[1]

Headline inflation, year on year, in per cent

Graph A.1

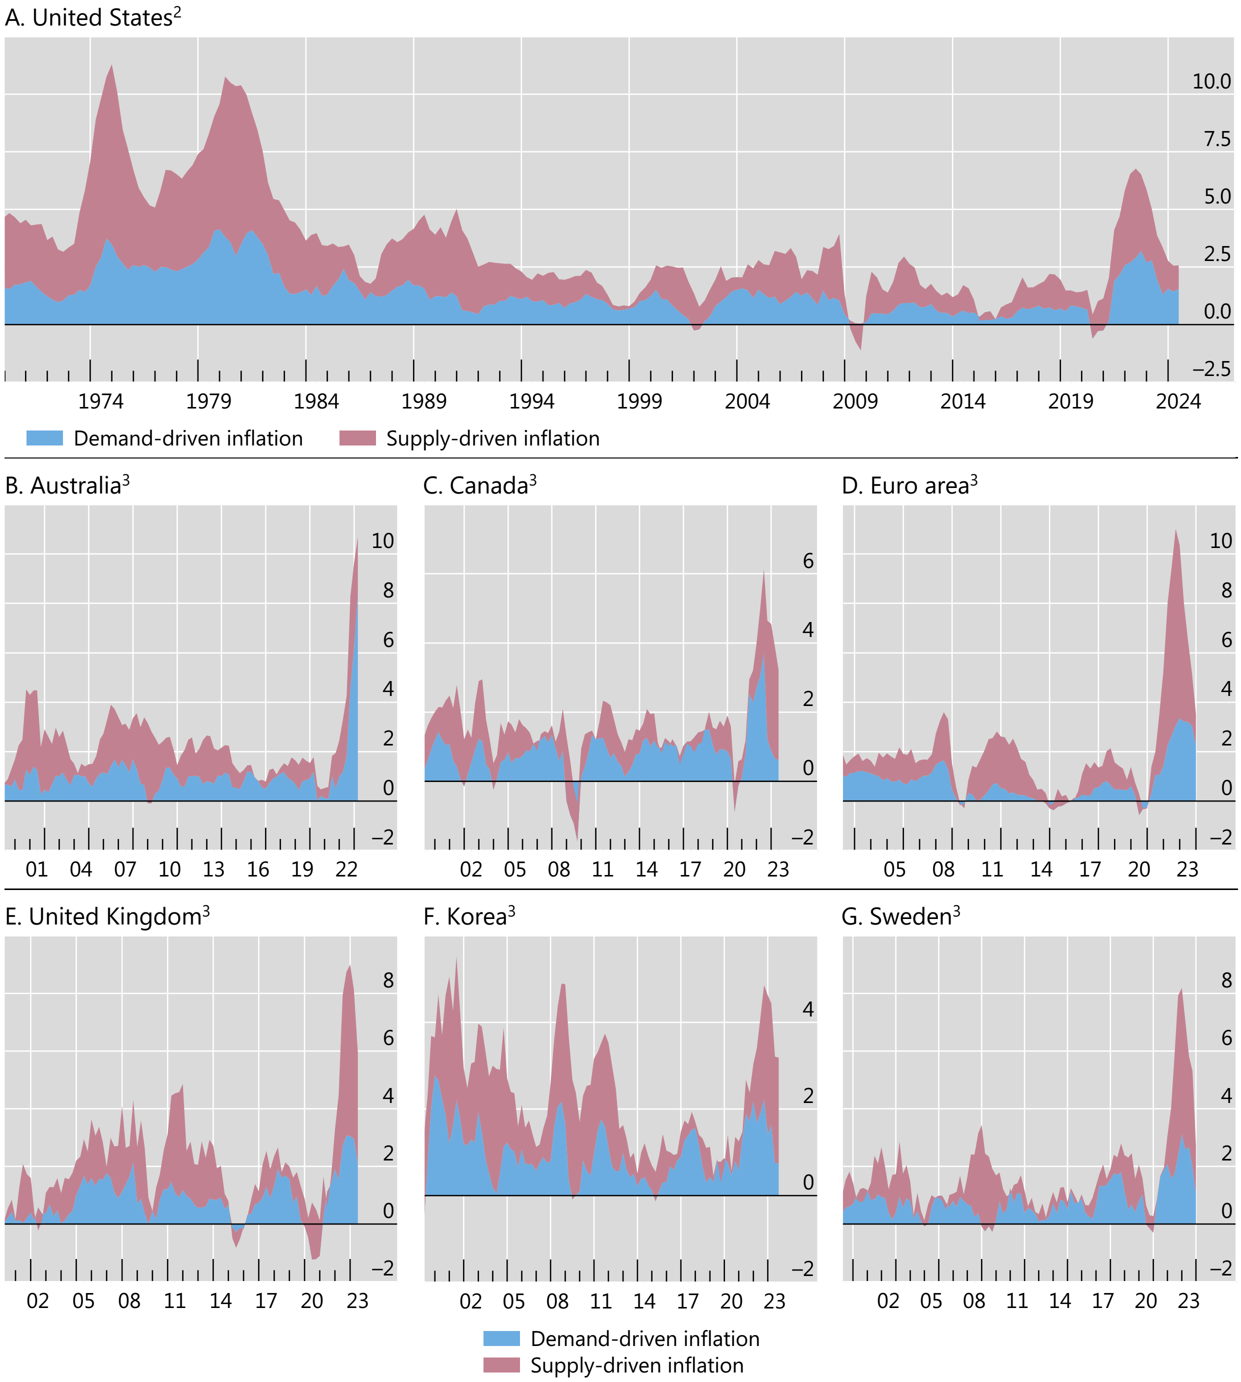


[1] The original Eickmeier and Hofmann (2022) inflation decomposition is expressed in quarter-on-quarter changes and concerns standardised inflation series. For our estimations, we use the mean and standard deviation of the aggregate inflation series to back out the demand and supply components of inflation, and then express those two components as year-on-year changes. [2] Decomposition based on the method in Shapiro (2022). Quarterly averages of monthly observations for the United States. [3] Decomposition based on the method in Eickmeier and Hofmann (2022).

Sources: Eickmeier and Hofmann (2022); Shapiro (2022); OECD; authors' calculations.

# Annex B

## Estimation procedure

Following the literature (eg Hofmann and Bogdanova (2012); Carvalho et al (2021)), we estimate the structural Taylor rule specifications in (1) and (2) based on reduced form OLS regressions. These regressions write:

- for the conventional Taylor rule:

$$i_t = \alpha + \rho^{aux} i_{t-1} + \alpha^{aux}\pi_t + \beta^{aux}\, y_t + \varepsilon_t$$

where the structural parameters of (1) are backed out from the auxiliary regression as follows: ρ= $\rho^{aux}$, α= $\alpha^{aux}$ /(1-$\rho^{aux}$), β= $\beta^{aux}$ /(1-$\rho^{aux}$).

- for the targeted Taylor rule:

$$i_t = \alpha^{aux} + \rho^{aux} i_{t-1} + \alpha^{d,aux}\pi_t^d + \alpha^{s,aux}\pi_t^s + \beta^{aux}\, y_t + \varepsilon_t$$

where the structural parameters of (2) are backed out as follows: ρ=$\rho^{aux}$, $\alpha^s$ = $\alpha^{s,aux}$/(1-$\rho^{aux}$), $\alpha^d$ = $\alpha^{d,aux}$/(1-$\rho^{aux}$), β = $\beta^{aux}$ /(1-$\rho^{aux}$).

## Estimation sample

The jurisdiction-specific samples are as follows: Australia: Q1 1999–Q1 2023; Canada: Q1 1999– Q2 2023; the euro area: Q1 2002 –Q4 2023; Korea: Q1 1999–Q3 2023; Sweden: Q1 1999–Q4 2023; the United Kingdom: Q1 2000–Q2 2023, the United States: Q3 1979 –Q2 2024 .

Byeungchun Kwon
*byeungchun.kwon@bis.org*

Taejin Park
*taejin.park@bis.org*

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

Phurichai Rungcharoenkitkul
*phurichai.rungcharoenkitkul@bis.org*

# Large language models: a primer for economists[1]

*Large language models (LLMs) are powerful tools for analysing textual data, with substantial untapped potential in economic and central banking applications. Vast archives of text, including policy statements, financial reports and news, offer rich opportunities for analysis. This special feature provides an accessible introduction to LLMs aimed at economists and offers applied researchers a practical walkthrough of their use. We provide a step-by-step guide on the use of LLMs covering data organisation, signal extraction, quantitative analysis and output evaluation. As an illustration, we apply the framework to analyse perceived drivers of stock market dynamics based on over 60,000 news articles between 2021 and 2023. While macroeconomic and monetary policy news are important, market sentiment also exerts substantial influence.*

*JEL classification: C55, C63, G10.*

Large language models (LLMs) represent a breakthrough application of machine learning techniques to natural language processing (NLP). Machine learning algorithms excel at imposing mathematical structure on unstructured data. They do so by converting text, speech or images into arrays of numbers, ie vectors. This "embedding" process has a wide variety of applications, as it transforms complex unstructured data into structured data suitable for mathematical operations and statistical analysis. LLM techniques can support many use cases, such as forecasting, nowcasting and surveillance, as well as sentiment analysis of news, social media and policy reports. For economists and central bankers accustomed to working with structured numerical data, LLMs are powerful additions to their toolkit.

This primer offers accessible guidance on LLM technologies and highlights key considerations for economists.[2] We cover topics such as model selection, pre-processing techniques, topic modelling, quantitative analysis and the involvement of human judgment. The guide is tailored to use cases commonly encountered by central bankers but with a wide applicability in any social science field that works with text data. To showcase the practical usage, we apply our guidelines to analyse perceived drivers of US equity prices, with accompanying sample code available in a GitHub repository. An online glossary provides definitions of key technical terms.

1 The views expressed are not necessarily those of the Bank for International Settlements. For helpful comments, we are grateful to Douglas Araujo, Claudio Borio, Leonardo Gambacorta, Gaston Gelos, Benoît Mojon, Andreas Schrimpf, Hyun Song Shin and Kevin Tracol. All remaining errors are ours.

2 Our focus on practical implementation complements recent reviews such as Araujo et al (2024), Athey and Imbens (2019), Dell (2024) and Korinek (2023, 2024).

*Key takeaways*

- *LLMs excel at organising text data into structured vector form. Unlocking their full potential requires careful planning, good research design and awareness of the tools' limitations.*
- *We discuss some best practices in deploying LLMs, including: (i) a modular workflow, (ii) informed choices of LLM tools and (iii) sufficient training data and examples for more complex tasks.*
- *Common pitfalls are: (i) unrealistic expectations of LLMs' capabilities, (ii) sub-optimal management of computational resources and (iii) insufficient use of human judgment in evaluating output.*

We first introduce key technologies that underpin LLMs. We then present a stylised workflow for their application, highlighting best practices and common pitfalls. Next, we show how to implement this workflow in a concrete exercise by isolating perceived drivers of US stock market movements. And finally we conclude.

## Introduction to large language models

### The central idea

Machine learning techniques excel at imposing mathematical structure on unstructured data. In the NLP context, this entails projecting words into a vector space – a process called *embedding*. Relationships between words are then represented by their Euclidean distance in the vector space, with a smaller distance indicating a closer semantic relationship. The word embedding for "football" is closer to "basketball" and further from "monsoon", which is closer to "cloud". These embeddings allow the use of algebraic techniques to convey relationships between words (see BIS (2024)). For example, the embeddings for countries and capitals would obey: $\overrightarrow{Seoul} = \overrightarrow{Korea} - \overrightarrow{Spain} + \overrightarrow{Madrid}$, consistent with the notion that Seoul is related to Korea in the same way that Madrid is related to Spain. Similarly, simple linear algebra applies, eg $\overrightarrow{Seoul} - \overrightarrow{Korea} = \overrightarrow{Madrid} - \overrightarrow{Spain}$. Furthermore, embeddings for sentences, paragraphs or any group of words can be computed as a weighted sum of the word embeddings to represent their collective meaning (Arora et al (2017)). Embeddings pave the way for applying mathematical tools to languages, enhancing tasks such as sentiment analysis, translation, sentence completion and named-entity recognition (eg recognising the Bank for International Settlements as a single unit rather than four separate words). This mapping of textual data into numerical form using embeddings is essential for any subsequent analysis.

Early neural networks[3] for NLP (eg Word2vec (Mikolov et al (2013)) assigned a single *unique* vector to each word. These methods, which were significant advancements at the time, have been extensively utilised by economists since their inception.[4] However, this one-to-one mapping suffers from an important drawback – it cannot recognise the *meanings of words that vary with context*. It might struggle with a sentence such as "The bank raises rates to lower inflation." Inferring correctly

[3] A neural network is a function that maps inputs to outputs through layers of matrix multiplications and a component-wise non-linear function that can provide a tight fit across a variety of tasks.

[4] See Matsui and Ferrara (2022) and the references therein.

that "bank" refers to the "central bank" rather than a riverbank or a commercial bank requires taking into account the words that give context in this sentence, namely "raises" and "rates".

What gives LLMs an edge over earlier NLP methods is a neural network known as the transformer architecture (Box A). This breakthrough creates a word embedding based on its context, allowing the embedding to capture the word's meaning within the overall context of the surrounding text rather than just its dictionary form.[5] For example, the transformer architecture would place embeddings for "bank" (financial) and "money" closer together because they frequently appear in similar contexts, while "bank" (river) would be positioned further away from "money" and closer to "meadow". Additionally, LLMs account for word sequences in sentences, enabling them to clearly distinguish between "The bank raises rates to lower inflation" and "The bank lowers rates to raise inflation."

## The techniques

The original transformer for translation included two components: an encoder, which transforms input language into an embedding vector, and a decoder, which converts the embedding into output language. Current LLMs draw on one of the two – the encoder as in bidirectional encoder representations from transformer (BERT) or the decoder as in generative pretrained transformer (GPT). Each has its own strengths. GPT *sequentially* uses each word's preceding text to create its contextualised embedding and can be used to generate text. This self-generating process is similar to how an autoregressive model is capable of making recursive forecasts. In contrast, BERT uses both preceding and subsequent words *taken together* to create the embedding for each word. It is akin to using the full sample to draw econometric inferences. GPT is more widely recognised but is not necessarily superior for all tasks. In many economic applications, BERT may be more suitable for quantitative analysis because it uses subsequent text to infer context (eg Gambacorta et al (2024)).

Economists can deploy LLMs to analyse vast amounts of text data more accurately and efficiently. LLMs are not a blank slate, as their parameters have been previously estimated with massive data sets downloaded from the internet — a process commonly known as *pretraining*. Economists can directly use these pretrained models to embed text for analysis. Alternatively, they can re-estimate or modify the LLM parameters using their economic text data, much like with any econometric model. This process, known as *fine-tuning*, adjusts the LLM to the specific economic data and research questions, yielding more accurate and relevant predictions.

Instead of fine-tuning, economists can also use the chatbot version of an LLM (eg ChatGPT, Claude, Gemini, etc) and ask the question directly through the chat interface or via an application programming interface (API). This approach does not modify model parameters but allows users to obtain improved responses by providing more context and guidance through prompting, known as *in-context learning* (Box B). Although this approach allows the convenient use of plain English or any other natural language, managing the consistency and quality of responses for robust research output can be challenging.

[5] Originally proposed for translation (Vaswani et al (2017)), the transformer removes bottlenecks in previous architectures by processing data in parallel, increasing speed by several orders of magnitude.

Box A

## Transformer models

Transformers are the core technology that underpins all current LLMs. They are parametric non-linear models built to generate embeddings of words while recognising their context. At the heart of the transformer architecture are two innovations: *multi-headed attention* and *positional encoding* (BIS (2024), Chapter III (Box B)). This box describes how these two mechanisms work together in transforming one-to-one word embeddings into contextualised vector embeddings (Graph A1).

The attention mechanism enables each word in a text to be interpreted in relation to all other words, enhancing the ability to consider context and relationships throughout the text. The mechanism is "multi-headed", as it uses several parallel attention layers to capture different meanings of the same word (purple layers in Graph A1). For instance, in the sentence "I sat on the south bank of the river Main overlooking the European Central Bank", each occurrence of the word "bank" will be mapped to a different vector embedding based on its surrounding words. The first will be closer to "water", and the second to "inflation". The attention mechanism quantifies the similarity between the word "bank" and all the words in the context, through score vectors or attention weights. Mathematically speaking, these are the dot products between the embedding vector of "bank" and key vectors representing each word in the text, including "bank" itself.

Positional encoding (orange blocks in Graph A1) enables transformers to process data concurrently rather than sequentially, easing a key constraint faced by earlier neural network models. By embedding each word with positional information, transformers preserve the sequence of words and enable parallelised training. This, in turn, allows training with more data and the building of bigger models, hence superior performance.

### Transformer model architecture[1]

Graph A1

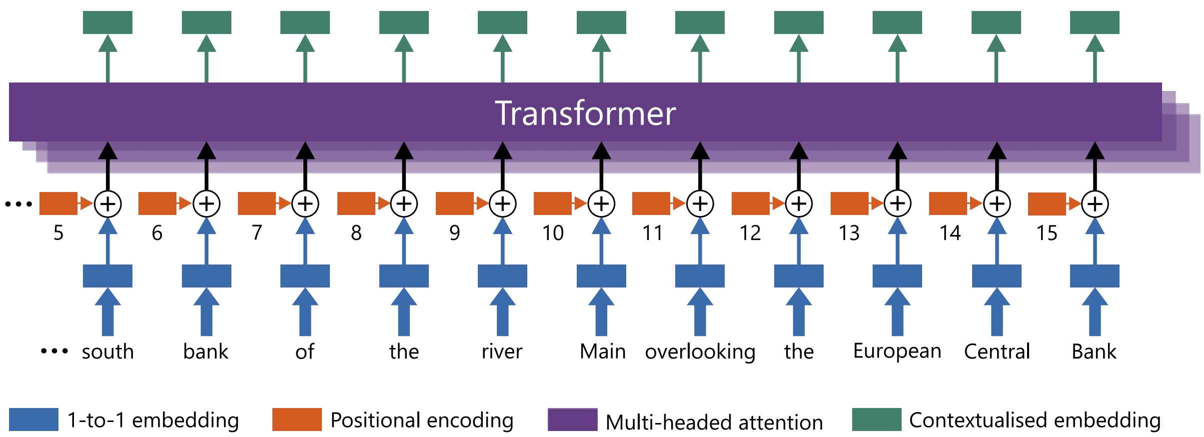


[1] Blue boxes represent the initial one-to-one word embeddings. Orange boxes are vectors of word positions in the paragraph, which are added to the one-to-one word embeddings before entering the transformer. Green boxes are the produced contextualised word embeddings. Purple boxes represent the transformer, made up of layers of multi-headed attention and fully connected neural networks. For example, GPT–3 (Brown et al (2020)) has 175 billion parameters, 12,288-dimensional word embeddings, 96 layers of multi-headed attention with 96 heads per layer, and 128-dimensional key and query vectors. See Vaswani et al (2017), Alammar (2018) and Cho et al (2024).

Source: Authors' elaboration.

Transformer models are characterised by four key features: the size of the training data, the size of the embedding vector, the context window (ie the length of text processed together in a batch) and the number of parameters. The quality of the embeddings usually improves with all four factors, with the size of the training data being the most important. The parameters are estimated using a vast amount of text data from sources such as the internet's Common Crawl, which includes books, web content, social media posts, news articles and more (totalling 570 GB, cleaned from 45 TB of original text) (Brown et al (2020)). This extensive training enables the transformer to map any sentence given its context into vectors. However, only a few organisations possess the computational resources and data to train these very large models from scratch. Consequently, most users apply open source LLMs or access proprietary models through APIs, referring to them as pretrained or foundational models.

Box B

## In-context learning and prompting

Chatbots based on LLMs can be used via APIs to directly ask questions – a method known as *in-context learning*, as it does not alter the LLM's parameters. For example, one can ask the chatbot to assess the sentiment of a sentence. This approach bypasses the usual empirical procedure in economics, which typically involves building statistical models, gathering data and manipulating them for diagnosis, causal inference or prediction. While the traditional empirical procedure often provides better predictions with sufficient data, LLM chatbots (eg ChatGPT, Claude, Gemini, etc) can yield surprisingly satisfactory results. This depends crucially, however, on how the question is framed – a process known as *prompting*. There are three main ways to prompt a chatbot.

The simplest way is to ask the question directly, eg: *Can you tell me if the following central banker statement is hawkish, neutral or dovish? "Recent data point to stabilisation, though latent risks tied to supply disruptions may warrant cautious optimism." Just indicate the sentiment to me without explanation*. This approach, known as *zero-shot* prompting, relies on the LLM's existing training. However, the results may not align well with research goals, especially when the task at hand is subjective and domain specific.

To improve accuracy, one can add examples to the prompt, eg: *I would classify [Sentence1] as hawkish, [Sentence2] as dovish, [Sentence3] as neutral,…*. This *few-shot* prompting method (Brown et al (2020)) enhances performance by guiding the model through examples. More examples improve LLM performance.

The third method explains the researcher's decision-making process to the chatbot itself, eg: *I would classify: [Sentence1] as hawkish because [Reason1], [Sentence2] as dovish because [Reason2], and [Sentence3] as neutral because [Reason3]*. This *chain-of-thought (CoT)* prompting has been known to outperform few-shot prompting in certain tasks, by guiding LLMs through a step-by-step reasoning process to help them produce more accurate responses (Wei et al (2022)).

Chatbot-based LLMs have several limitations in research applications. Once the context is reset, the chatbot forgets previous questions and examples. It is also important to emphasise that LLM-based chatbots predict statistically plausible outcomes but cannot ensure factual accuracy. Human oversight is therefore critical to interpret and validate results. And because changes in the prompt can alter outcomes, researchers should test multiple prompts using validation data wherever possible.

## Putting LLMs to work

To show how to put LLMs to work, we lay out a step-by-step workflow analogous to that of an econometrician and discuss how LLM tools can enhance capabilities in analysing unstructured text data at scale.[6] The workflow includes the following steps:

1. **Data organisation:** Just as an economist begins an empirical project by collecting and cleaning data, an analysis of text starts with text retrieval, pre-processing and, additionally, vector embedding.
2. **Signal extraction**: The next step is to extract key informational content from the data, similar to principal component analysis (PCA) or trend filtering in econometrics. This is important because the embedding vectors from the previous step have many dimensions (eg 4,096 for Llama 3.1 8B) and are too complex to analyse. Hence, some dimensionality and complexity reduction are necessary. Topic modelling is a powerful tool to do this, by organising text into common themes while respecting the context.

[6] Researchers could readily extend this analysis to images and audio too.

3. **Quantitative analysis**: In this step, the researcher engages in quantitative modelling and analysis. Sentiment analysis is one common tool for measuring intent, tone and opinions from text and can be assisted by LLMs to identify meaningful labels or scores. Standard statistical methods such as regression may also be applied directly to embeddings or sentiment scores produced by LLMs.
4. **Outcome evaluation**: The last step is to evaluate the model's output, eg based on the out-of-sample prediction as in standard econometrics. Here, text data raise unique challenges.

We consider each of the above steps in more detail.

## Data organisation

The objective of this step is to gather relevant text data and prepare them for contextualised embedding through text cleaning. The pertinent sources may be central bank statements, governors' speeches, policy papers, news articles or any other text sources. The next step is to divide the documents into so-called "chunks" of words, ie text segments of words, to prepare for embeddings (step 2 of Graph 1).[7] Each chunk serves as a unit of text analysis. The chunk size strikes a balance between breadth and granularity. It should be long enough to capture the relevant information and context, but not so long as to collapse a large block of text into a single vector and lose the nuances of different messages.

The next step involves breaking words down to their root forms to simplify the task of LLM embeddings. *Tokenisation* is a process of breaking words down into smaller units called tokens; *lemmatisation* in turn simplifies words to their base forms to focus on their core meanings (step 3 of Graph 1). For example, tokenisation breaks the word "disinflationary" down into three tokens: "dis-" (prefix), "inflation" (base word) and "-ary" (suffix). Lemmatisation assigns "lowered" to "lower", "pressures" to "pressure" and "tracks" to "track". Each LLM has its own set of rules for tokenising and lemmatising words.[8] Researchers need to apply appropriate tokenisation rules that are compatible with their selected LLM.

Once all text chunks have been tokenised and lemmatised, one can proceed to generating the embeddings (step 4 of Graph 1). If using a BERT-like (encoder) model, one will generate a global embedding for the entire chunk, called the classification (CLS) embedding. If using a GPT-based (decoder) model, the embedding for the text chunk can be obtained as a weighted sum of each word's embedding (Arora et al (2017)). In either case, one can rely on embedding models that have been trained by tech companies such as Anthropic, Google, Meta, Mistral AI or OpenAI. Alternatively, one could modify these models using own data[9] to tailor the embeddings to research objectives. This *domain-adaptation* option is naturally more costly and is only advisable if sufficient data and computational resources are available. Additionally,

7 Dividing text data into blocks of words – eg several hundred words as a proxy for a paragraph – is often easier to program and hence more practical than differentiating sentences or paragraphs.

8 For examples, see the OpenAI Platform Tokenizer at platform.openai.com/tokenizer or the Lunary Llama 3 Tokenizer at lunary.ai/llama3-tokenizer.

9 Training these models with private data is not advisable if the models are available for widespread use, as they may memorise the training data.

## Stylised text data processing flow[1]

Illustration of topic modelling

Graph 1

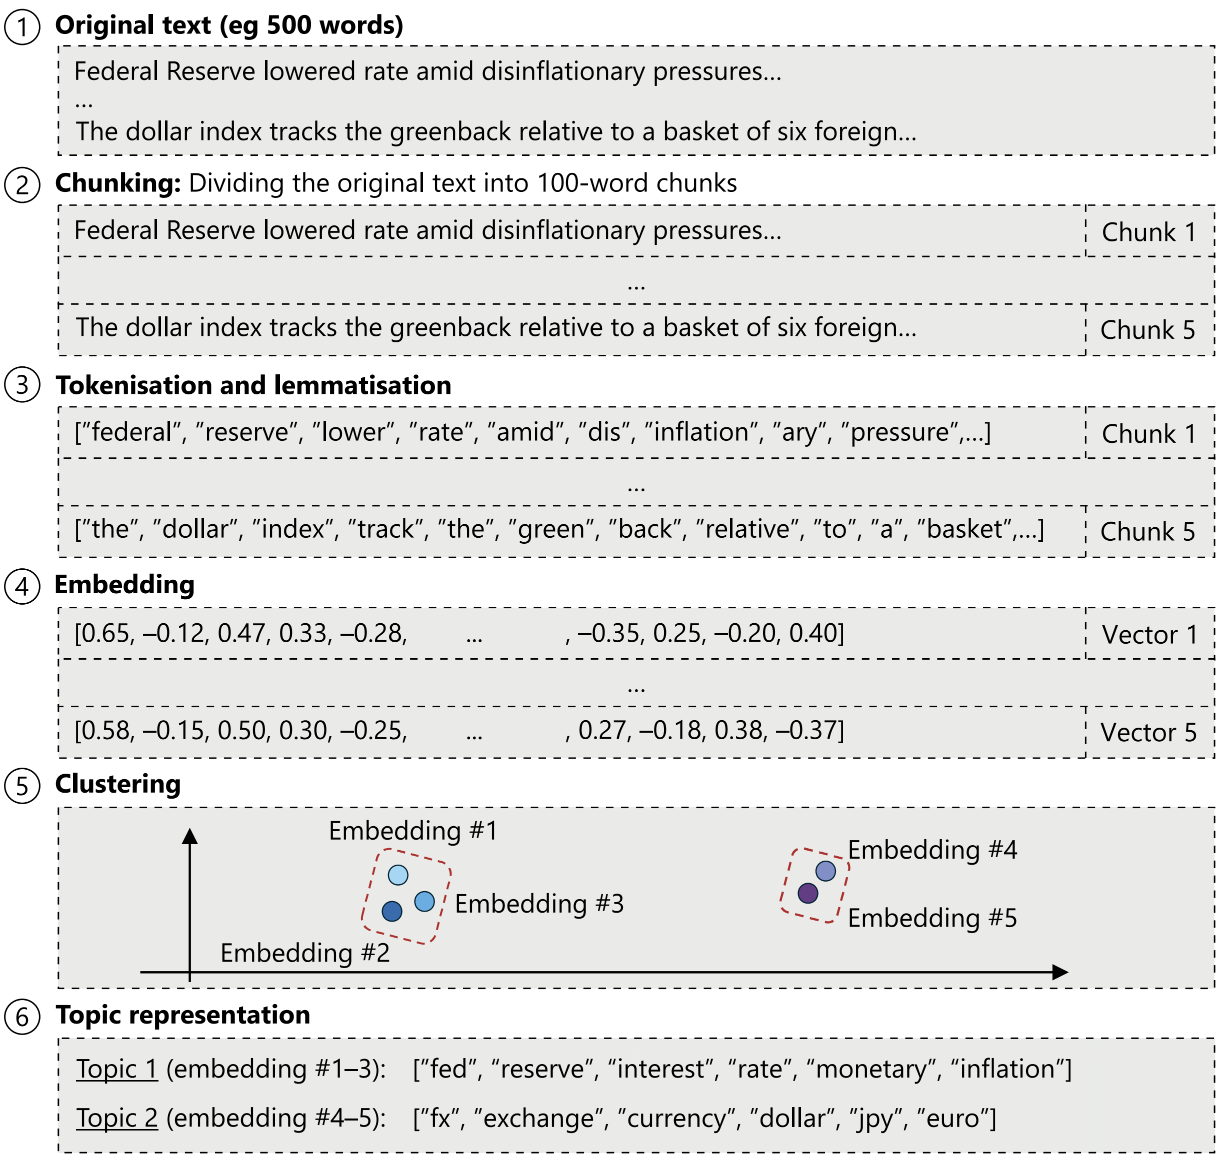


[1] In Step 4, only the aggregated embeddings at the chunk level are shown, although they are initially produced at the token level.

Source: Authors' elaboration.

the recent state-of-the-art LLMs are often sufficiently trained with economic terms, lessening the need for additional training.

## Signal extraction

The second step aims to extract signals and distil key information from the retrieved text data through dimensionality reduction. It simplifies high-dimensional embedding vectors by extracting only the most relevant information needed for meaningful interpretation, thus improving computational efficiency. In LLM analysis, potential tools include topic modelling, clustering techniques or embedding-based dimensionality reduction, such as Uniform Manifold Approximation and Projection (UMAP) (McInnes et al (2018)). These methods facilitate the formation of clusters of embeddings obtained from the previous step (step 5 of Graph 1).

We focus on **topic modelling** due to its broad applicability, where the objective is typically to extract the main messages from text data. Traditional topic modelling tools, such as latent Dirichlet allocation (LDA) (Blei et al (2003)), do this by grouping keywords under common themes but cannot differentiate between nuanced subtopics or contextual shifts within the same overarching theme. With LLMs, topic modelling leverages the transformer architecture and embeddings to provide a context-sensitive representation of text in identifying the dominant themes.

Topic modelling can be *unsupervised*, allowing the data to reveal patterns autonomously, or *supervised*, where prior knowledge or guidance shapes the output.

*Unsupervised* topic modelling summarises the dominant themes in text through keyword lists, offering a high-level overview of key topics without relying on pre-defined classification schemes. A common procedure, such as BERTopic (Grootendorst (2022)), starts by generating embeddings for each text chunk. The algorithm then groups contextually similar embeddings into clusters, typically after applying dimensionality reduction to enhance efficiency and performance (see step 5 of Graph 1). After completing this step, a researcher may assign a topic label – a few keywords – that best represents the cluster. These labels are often selected manually from the original list of keywords associated with the cluster, both reflecting the researcher's focus and interests and ensuring clear differentiation between clusters. For example, in step 6 of Graph 1, one might label Topic 1 as "monetary policy" and Topic 2 as "FX markets".

A potential drawback of unsupervised topic modelling is that it may generate topics irrelevant to the researcher's focus. For example, when analysing stock markets, names of major stocks or asset management firms might dominate because they are frequent within a specific cluster and distinct from those in other clusters. This may not be of interest to a macro-finance researcher who wants to study monetary policy and market valuations. As a result, unsupervised methods often require post-processing to align the topics with specific research goals – a practice we follow in the application below.

Resorting to *supervised* topic modelling avoids generating irrelevant topics, as it allows researchers to define topics of interest from the outset. One common option is *seeding*, where a set of relevant terms is introduced to guide the topic modelling process. For example, providing keywords related to monetary policy (eg "Federal Reserve", "policy rates", "FOMC Chair") directs the model to focus on text segments that are semantically closest to the concept of monetary policy.

## Quantitative analysis

Once the text data have been converted into vectors and organised by topics, they are ready for use in quantitative analysis. As the embeddings are in vector format, researchers can employ standard econometric tools such as panel regressions for classification data or autoregressions for time-stamped text. In the illustration that follows, we focus on **sentiment analysis**. This is one of the most relevant problems in economics and social science when the data are only in text form.

Sentiment analysis assigns predefined labels to each text chunk. Classification could be binary (positive or negative), tertiary (positive, neutral or negative) or more multidimensional.[10] Without an LLM, this transformation may need to rely on a dictionary of terms compiled by the researcher to map a chosen list of words reflecting tones to sentiments. This approach is laborious and prone to errors and risks false assignments when context is important.

Depending on the application, one can analyse sentiment either at the small chunk level (local analysis) or globally across the entire document. The chunk size sets the level of granularity: smaller chunks allow for different passages in an article to convey different sentiments. In some cases, capturing local sentiment may be more advisable, as the goal may be to uncover different aspects of the unfolding of a story. For instance, both positive and negative forces weighing on the stock market could be described in the same document. In other cases, the document's overall message may be where the researcher's interest lies. For instance, in monetary policy statements, the overall assessment and final decision may matter the most, justifying document-level sentiment analysis.

The choice of language models for sentiment analysis depends on many factors. These include task complexity, text volume, the availability of labelled training data, computing capacity, expertise and the time required for manual labelling. Among sentiment analysis tools, pretrained encoder-based models such as BERT (Devlin et al (2018)) and RoBERTa (Liu et al (2019)) offer a suitable architecture for most applications. These models are specifically designed to encode contextual information, excel at capturing the sentiments of input text and are fast to run. For highly specialised applications, specifically trained LLMs can be more effective. For instance, the "CB-LMs" of Gambacorta et al (2024) are trained on central banking documents and may be better at capturing central banking nuances than general models. As for decoder-based GPT models, their extremely large sizes are key advantages, though their architecture may not be ideal for sentiment classification and they tend to have a longer runtime.

If the researcher opts for BERT-based LLMs, they will need to conduct fine-tuning to help the model learn how to classify embeddings. This requires a labelled data set (eg assigning positive, neutral and negative to a portion of the data as training examples). During fine-tuning, the model is further trained on this labelled data set in a supervised learning process to classify sentiments in the embedding space (Devlin et al (2018)). The process updates both the embeddings themselves and the boundaries separating different sentiments, making it easier for the model to correctly classify sentiments of new text inputs.

If the researcher adopts GPT-based models which have been trained with a large amount of data, such as state-of-the-art chatbots, fine-tuning is often unnecessary and not worthwhile.[11] In many cases, low-effort in-context learning can be sufficient (see Box B and Gambacorta et al (2024)). These models typically require minimal or no data preprocessing, such as tokenisation or embedding. Researchers can directly input text for topic modelling and sentiment analysis using common language instructions.

[10] One example is to capture the intensity and nuances of opinions expressed in economic or financial discourse. For instance, in analysing central bank communication, one can detect degrees of confidence and uncertainty about the economic outlook beyond binary classification.

[11] Fine-tuning GPT-based models is computationally demanding and generally does not improve performance over fine-tuning smaller BERT-based models for simple classification tasks.

## Outcome evaluation

LLMs, as a machine learning tool, are best evaluated by their predictive performance. For example, in classification problems, the objective could be to minimise the misclassification error. Note that usual statistical principles such as parsimony or statistical significance of coefficients do not apply in a machine learning context. Modern LLMs do in fact have hundreds of billions of parameters and embedding vectors with a few thousand elements. The main focus is on making accurate predictions in the test set, with overfitting concerns addressed by various techniques (eg regularisation, initialisation and stochastic gradient descent).

In evaluating performance, researchers should choose error measures that align with their applications and desired outcomes. For example, in regression with outlier problems, absolute errors might be preferable to squared errors. Or when the outcome spans several orders of magnitude, measuring errors in the logarithm of the output may help reduce the relative error. For problems with imbalanced classes and uneven costs for false positives and false negatives, the F1 score[12] is often an appropriate metric, as it trades off "precision" and "recall" in classification. Precision measures how many predicted positive examples are positive, while recall captures how many positive examples are correctly identified by the model. Combining the two helps control both false positives and false negatives.

# Illustration in a study of equity market drivers

We now illustrate how to operate the workflow in a concrete project. The objective is to identify perceived drivers of stock market prices using news reports. The sample includes 63,388 daily news articles published over 2021–23, from sources such as the *Wall Street Journal*, Reuters, *Forbes* and MarketWatch. The sample period covers a sharp monetary policy tightening and a surge in inflation. Graph 2 summarises the overall workflow.

### Data organisation

We start by retrieving news articles from Factiva and use its topic tagging to select articles potentially relevant to the US stock markets, as measured by their performance in the S&P 500. We split each of the 63,388 articles into chunks of 100 words to capture localised context, producing 349,896 text chunks.[13] We transform each chunk using the Llama 3.1 8B, which embeds each word in the chunk into a 4,096-dimensional vector. To summarise the key messages of each chunk, we then average over the 100-word embeddings, to obtain a 4,096-dimensional vector embedding representing the whole chunk.[14]

[12] The F1 score is the harmonic mean of precision and recall: $\frac{2 \times Precision \times Recall}{Precision + Recall}$, where Precision = True Positives / All Positives, and Recall = True Positives / (True Positives + False Negatives).

[13] The chunk size should reflect research objectives, eg in our case, to differentiate topics roughly corresponding to paragraphs. While a 100-word chunk is standard, this can vary with applications.

[14] Chunk embeddings can also be obtained through proprietary API services, eg as offered by OpenAI.

Workflow for analysing stock market drivers using news articles Graph 2

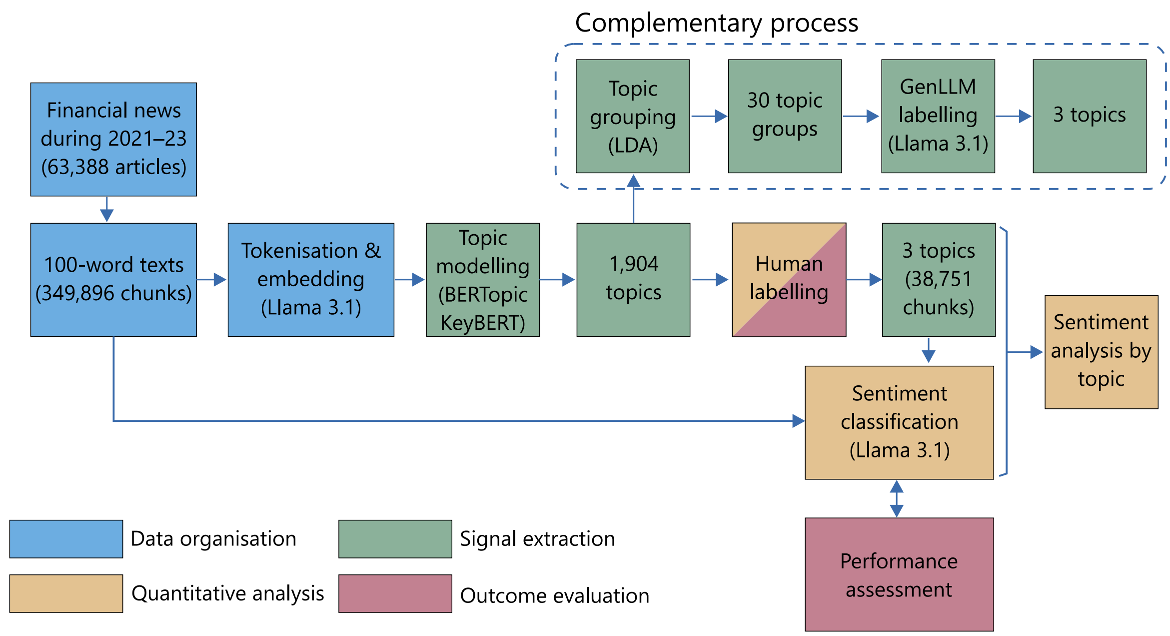


Source: Authors' elaboration.

## Signal extraction

We next employ BERTopic to assign text chunk embeddings to initial topics. As the initial output contains several irrelevant topics, some filtering is necessary. A significant portion of the chunks (190,477) are dropped, as they do not relate to any specific topic clusters[15] and instead cover general topics (eg economy, policy, market, etc). Additionally, we filter out chunks that belong to a cluster but are only loosely associated. We also exclude clusters with too few articles.

To further refine the topic sets, as illustrated in Table 1, we apply KeyBERT (Grootendorst (2020)) to extract meaningful keyword sets from the BERTopic results and filter out less relevant terms. We then merge the refined keyword sets using LDA,[16] reducing the nearly 2,000 topics into a smaller, more manageable set of keyword clusters. LDA allows us to analyse the topics within broader themes.[17]

To assign intuitive and descriptive labels to the topics, we use Llama 3.1 to generate initial labels based on the grouped keyword sets. One could also apply

[15] We employ a hierarchical density-based clustering method (HDBSCAN) for clustering. It groups similar items by finding densely packed areas in the data, forming clusters within clusters to reveal patterns at different levels of detail.

[16] LDA is a Bayesian model that identifies topics within a text corpus by assuming that documents are mixtures of topics and that topics are mixtures of words. Each document is represented as a mixture of topics, with each word assigned to a particular topic.

[17] It is possible to minimise the number of topics directly from BERTopic. However, starting with more detailed topics from BERTopic, and then using LDA to consolidate them later in a modular way, offers better control over the process and provides greater transparency and clarity for the analysis.

| Example of outputs in each step of the topic modelling process | Table 1 |
|---|---|
| Topic modelling step | Output |
| Initial topic representation (eg BERTopic) | ["the", "of", "in", "to", "company", "business", "profit", "revenue", "estimate", "price", "mix", "for", "year", "quarter", "earnings", "cent", "margin", "store", "increase", "loss", "delta", "adjusted", "per", "eps", "capacity", "result", "sale", "million", "a", "by", "analyst", "income", "expense", "fiscal"] |
| Fine-tuned representations (eg KeyBERT) | ["revenue", "profit", "income", "estimate", "company", "expense", "fiscal", "quarter", "earning", "sale", "forecast", "margin", "increase", "loss", "adjusted", "delta", "result", "capacity", "year", "eps"] |
| Topic grouping (eg LDA) | ["company", "earning", "estimate", "forecast", "profit", "quarter"] |
| GenAI label (eg Llama 3.1) | Equities and corporate earnings |
| Expert label | Fundamentals |

Sources: Dow Jones Factiva; authors' calculations.

Llama 3.1 from the initial step of topic modelling. That said, as discussed, the large input size and computational overhead pose a constraint. Additionally, maintaining control over the process ensures more effective and transparent topic labelling.

Finally, we review the topics and make final manual adjustments to ensure the labels and groupings are contextually meaningful and aligned with research objectives. The initial topic grouping performed by LDA, along with the labelling provided by Llama 3.1, facilitates this process. After reviewing, we categorise the news articles into three broad topics: *fundamentals* (with KeyBERT keywords such as "profit", "economy" and "employment"), *monetary policy* ("fed" and "interest rate"), and *market sentiment* ("overbought", "buzz", "volatility" and "ipos").

This classification into three topics mirrors common analytical approaches to stock price analysis. First, stock prices are determined by current and expected future dividends (fundamentals). These are adjusted by the discount rate, which depends on the risk-free interest rate (monetary policy) as well as a compensation for risk that reflects investor risk appetite and sentiment (market sentiment).

Topic modelling reduces the data set to just 10% of its original size, resulting in a final output of 38,751 targeted text chunks, assigned to the three categories: fundamentals (16,879), market sentiment (14,468) and monetary policy (7,404).

## Quantitative analysis and outcome evaluation

After classifying news article texts into the three topics, we conduct sentiment analysis within each topic. We use the Llama 3.1 70B model[18] to classify each topic in each news chunk as either negative, neutral or positive (with scores –1, 0 and 1), with positive indicating an association with stock price increases. In doing so, we employ few-shot learning and provide task-specific instructions and examples for each

[18] Although BERT-based models are often suitable for sentiment classification tasks, two key factors motivated our choice of a GPT-based model. First, the length of each chunk (100 words) makes the classification task more complex, requiring a high-performance model. Second, topic modelling has reduced the number of articles to a size that is feasible for processing by a large model.

topic.[19] We then sum the scores across all article chunks to derive the aggregate sentiment score corresponding to each topic.

### Results

The aggregate sentiment score aligns well with stock market movements (Graph 3.A). This suggests that the procedure is extracting the appropriate signals from text data. Examining the relative importance of sentiment scores across the three categories, "market sentiment" and "fundamentals" co-move the most with the stock returns, showing correlation coefficients of 0.64 and 0.52, respectively (Graph 3.B).

The association between monetary policy sentiment and stock market performance is more tenuous, with a correlation of 0.30. Despite this period being one of the most aggressive monetary tightening episodes in recent history, the stock market – aside from the large drawdown in 2022 – held up well. One possible interpretation is that strong "fundamentals", and hence associated positive "market sentiment", may have offset the dampening effect of "monetary policy" tightening on stock returns.

Over the sample, the relationship between monetary policy sentiment and stock market performance has seen notable fluctuations, unlike market sentiment and

Factors associated with US stock market movements as reflected in news text

During the period of monetary policy tightening

Graph 3

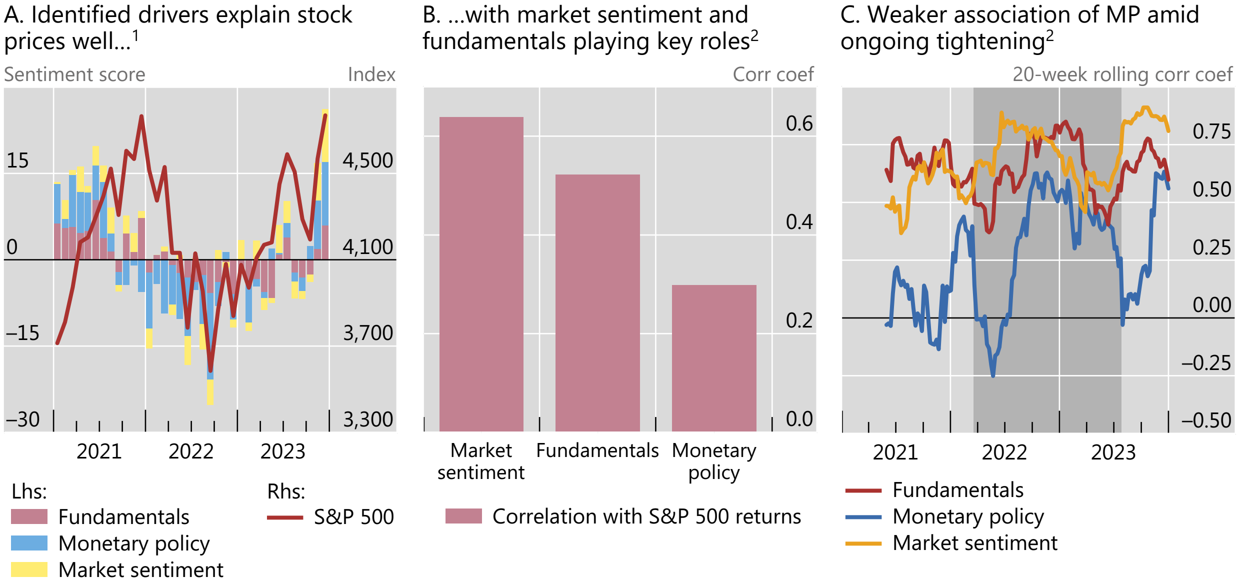


The shaded area indicates 17 March 2022–26 July 2023 (period of monetary policy tightening).

[1] Daily sentiment scores (–1: negative; 0: neutral; +1: positive) are averaged by category and aggregated monthly. [2] Data reflect weekly sentiment score sums and S&P 500 weekly returns between 2021 and 2023.

Sources: Bloomberg; Dow Jones Factiva; authors' calculations.

[19] For example, the prompt for fundamentals includes "Assign a sentiment score of –1 for negative, 0 for neutral, or 1 for positive. A positive sentiment means earnings or their outlook are upbeat, GDP growth prospects are bright, labour market remains tight [...]. A negative sentiment is the opposite [...]. Neutral sentiment is when the fundamental developments have no clear implications [...]."

fundamental factors. The correlation between monetary policy sentiment and stock returns rose to its highest point of over 0.6 in 2022, coinciding with the most active monetary policy tightening. But once the tightening stopped, the correlation between monetary policy sentiment and stock returns declined (Graph 3.C). The result is consistent with the broader disconnect between financial conditions and the monetary policy stance in recent years, as documented elsewhere. [20]

## Final considerations

Our stylised workflow offers broad lessons on general principles and best practices in making the most of LLMs. One is careful resource planning. While extremely powerful, the use of LLMs can be computationally expensive. Some of the computational steps, such as working with high-dimensional embeddings or the largest LLMs, may be run at most a few times in practice. This also means good research design is critical. For example, condensing input text data, as we do using topic modelling, helps reduce unnecessary computational load and improve model performance by focusing resources on the most relevant tasks. More generally, adopting a modular step-by-step process would enable researchers to check and analyse intermediate outputs, ensuring that the LLM procedure produces sensible outcomes as intended. Finally, given the probabilistic nature of LLM algorithms, some repetition of the procedures would help ensure robustness, if at additional computational costs.

Another important consideration pertains to data privacy and confidentiality. Researchers must ensure compliance with relevant regulations, such as corporate policies, data licences or other data protection laws. Using locally hosted LLMs, instead of cloud-based ones, may mitigate risks to some extent but not entirely, as some regulations apply regardless of the hosting environment.

The vast power of LLMs is best harnessed when researchers both understand and manage the models' limitations effectively. Recognising and addressing potential biases is essential at each stage of the analysis to ensure a successful application. As LLMs are in essence statistical models, it is up to users to apply informed judgment in evaluating outputs. The synergy between human oversight and LLM capabilities can help enhance quality control while mitigating potential risks, leading to more robust outcomes.

[20] See *BIS Quarterly Review* (March 2024) for example.

## References


Alammar, J (2018): "The illustrated transformer", blog post, available at jalammar.github.io/illustrated-transformer/.

Araujo, D, S Doerr, L Gambacorta and B Tissot (2024): "Artificial intelligence in central banking", *BIS Bulletin*, no 84, January.

Arora, S, Y Liang and T Ma (2017): "A simple but tough-to-beat baseline for sentence embeddings", in *Proceedings of the 5th International Conference on Learning Representations*, Toulon, France, 24–26 April.

Athey, S and G W Imbens (2019): "Machine learning methods that economists should know about", *Annual Review of Economics*, vol 11, pp 685–725, August.

Bank for International Settlements (2024): *Annual Economic Report,* June.

Blei, D, A Ng and M Jordan (2003): "Latent Dirichlet allocation", *Journal of Machine Learning Research*, vol 3, pp 993–1022.

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, D Ziegler, J Wu, C Winter, C Hesse, M Chen, E Sigler, M Litwin, S Gray, B Chess, J Clark, C Berner, S McCandlish, A Radford, I Sutskever and D Amodei (2020): "Language models are few-shot learners", *arXiv*, no 2005.14165 [cs.CL].

Cho, A, G Kim, A Karpekov, A Helbling, J Wang, S Lee, B Hoover and P Chau (2024): "Transformer explainer", *arXiv*, no 2408.04619, available at poloclub.github.io/transformer-explainer/.

Dell, M (2024): "Deep learning for economists", *arXiv*, no 2407.15339v1 [econ.GN], July.

Devlin, J, M-W Chang, K Lee and K Toutanova (2018): "BERT: pre-training of deep bidirectional transformers for language understanding", *arXiv*, no 1810.04805 [cs.CL].

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.

Grootendorst, M (2020): "KeyBERT: minimal keyword extraction with BERT", Zenodo, version v0.3.0, doi.org/10.5281/zenodo.4461265.

Grootendorst, M (2022): "BERTopic: neural topic modeling with a class-based TF-IDF procedure", *arXiv*, no 2203.05794 [cs.CL].

Korinek, A (2023): "Generative AI for economic research: use cases and implications for economists", *Journal of Economic Literature*, vol 61, no 4, pp 1281–1317, December.

Korinek, A (2024): "LLMs level up—better, faster, cheaper: June 2024 update to Section 3 of 'Generative AI for Economic Research: Use Cases and Implications for Economists'", published in the *Journal of Economic Literature*, vol 61, no 4.

Liu, Y, M Ott, N Goyal, J Du, M Joshi, D Chen, O Levy, M Lewis, L Zettlemoyer and V Stoyanov (2019): "RoBERTa: a robustly optimized BERT pretraining approach", *arXiv*, no 1907.11692 [cs.CL].

Matsui, A and E Ferrara (2022): "Word embedding for social sciences: an interdisciplinary survey", *arXiv*, no 2207.03086 [cs.AI].

McInnes, L, J Healy and J Melville (2018): "UMAP: uniform manifold approximation and projection for dimension reduction", *arXiv,* no 1802.03426 [stat.ML].

Mikolov, T, K Chen, G Corrado and J Dean (2013): "Efficient Estimation of Word Representations in Vector Space", *arXiv,* no 1301.3781 [cs.CL].

Vaswani, A, N Shazeer, N Parmar, J Uszkoreit, L Jones, A Gomez, L Kaiser and I Polosukhin (2017): "Attention is all you need", *arXiv*, no 1706.03762 [cs.CL].

Wei, J, X Wang, D Schuurmans, M Bosma, V Chi, Q Le, E Chi and D Zhou (2022): "Chain-of-thought prompting elicits reasoning in large language models", *arXiv*, no 2201.11903 [cs.CL].

Ryan Banerjee
*ryan.banerjee@bis.org*

Denis Gorea
*denis.gorea@bis.org*

Deniz Igan
*deniz.igan@bis.org*

Gabor Pinter
*gabor.pinter@bis.org*

# Monetary policy and housing markets: insights using a novel measure of housing supply elasticity[1]

*This special feature explores how the responsiveness of new housing supply to house prices (the supply elasticity) affects the transmission of monetary policy. We document a long-term decline in the supply elasticity of new housing and find that a less elastic supply amplifies the impact of monetary policy on house prices compared with rents. This disparity could raise questions about the measurement of housing costs in inflation measures targeted by central banks, as typically rents, rather than house prices, directly enter consumer price indices. Our findings highlight the importance of considering changing housing supply conditions in monetary policy formulation.*

*JEL classification: E52, R31.*

There have been strong secular trends in housing supply over the past few decades. In advanced economies (AEs), housing starts have been on a steady downward trend since the 1970s (Graph 1.A), accompanied by a rise in house prices relative to rents (Graph 1.B). This contrasts with the general trend in emerging market economies (EMEs) where housing starts have steadily ramped up since the 1990s, while the ratio of house prices to rents has remained more stable. These differences in housing supply conditions appear to play a crucial role in shaping the relationship between house prices and rents.

Against this backdrop, we explore how the responsiveness of new housing supply to price changes (the supply elasticity) influences the transmission of monetary policy to housing markets. As a first step, we develop time-varying measures of the supply elasticity for 21 economies, with data dating back to the 1970s in some cases. We approximate the supply of new housing with construction permits. We then analyse how the supply elasticity affects the transmission of monetary policy to house prices and rents. By incorporating time-varying elasticities, our analysis complements existing studies, which have relied on largely unchanging factors, such as geographical or regulatory constraints, to capture the elasticity. And by considering

[1] The views expressed are not necessarily those of the Bank for International Settlements. We thank Burcu Erik for excellent research assistance. For helpful comments, we are also grateful to Frederic Boissay, Claudio Borio, Mathias Drehmann, Gaston Gelos, Boris Hofmann, Enisse Kharroubi, Emanuel Kohlscheen, Benoît Mojon, Daniel Rees, Tom Rosewall, Andreas Schrimpf and Hyun Song Shin. All remaining errors are ours.

*Key takeaways*

- *The responsiveness of new housing supply to changes in house prices (the supply elasticity) has declined over the past five decades.*
- *When housing supply is inelastic, a given interest rate change leads to greater adjustments in the price-to-rent ratio, almost entirely accounted for by the larger impact on house prices. Rents barely respond irrespective of the supply elasticity.*
- *Larger house price responses due to a declining supply elasticity could strengthen the monetary transmission via its effect on homeowners' balance sheets. The limited traction on rents indicates that a considerable portion of the consumer price index is unresponsive to monetary policy.*

cross-country differences, our analysis extends existing studies, which focus on experience in the United States or only on house prices.[2]

Our results show that the supply elasticity influences the transmission of monetary policy to housing markets. When supply is less elastic, monetary policy easing has a larger impact on the price-to-rent ratio. This difference is almost entirely accounted for by the larger impact on house prices. Rents, by contrast, are much more sticky, irrespective of supply elasticity.

What could explain the findings? As in any other market, when supply is inelastic, more of the market adjustment to a change in demand – here monetary policy-induced – takes place through prices. This, in turn, could amplify the house price movements to the extent that it encourages households to extrapolate current house price growth into the future.

Divergent trends in housing markets across country groups[1] Graph 1

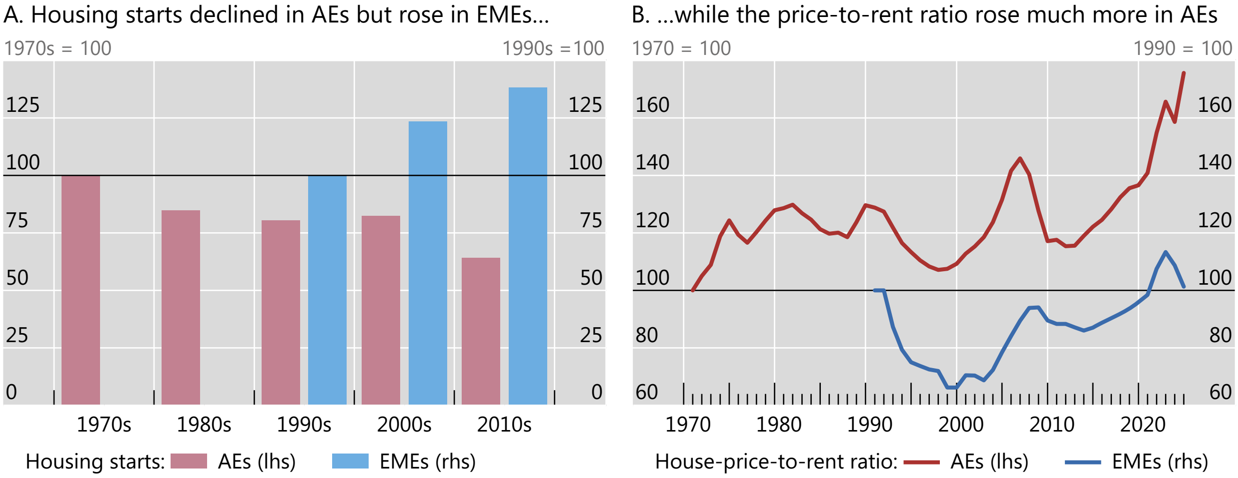


[1] See technical annex for details.

Sources: OECD; LSEG Datastream; Macrobond; national data; authors' calculations.

[2] For example, Dias and Duarte (2019) study the effect of monetary policy shocks on rents in the United States, while Aastveit and Anundsen (2022) and Albuquerque et al (2024) assess how geographical and regulatory differences across US states influence the monetary transmission to house prices and rents, respectively. Andaloussi et al (2024) examine how cross-country differences in population density influences the transmission of monetary policy to house prices.

These findings have implications for policy. Where housing supply has become less elastic, the larger impact of monetary policy on house prices may strengthen the transmission through homeowners. As shown by Cloyne et al (2020), house price swings amplify especially the consumption response of mortgagors. Another implication arises from the larger disparity between the monetary impact on house prices and rents when housing supply is inelastic. As rents, rather than house prices, enter consumer price indices (CPIs) directly, this could raise questions about the measurement of housing costs in inflation measures targeted by central banks.[3]

This special feature proceeds as follows. The first section discusses our newly proposed measurement of housing supply elasticity. The second section analyses how the supply elasticity influences the transmission of monetary policy to housing markets and discusses potential mechanisms. The final section draws some policy implications.

## Measuring housing supply elasticity

We estimate housing supply elasticity and how it changes over time using novel methods. Supply elasticity is a measure of how responsive housing supply is to changes in house prices. As it is not directly observable, it needs to be estimated. We estimate it using novel econometric techniques applied to data on house prices, building permits (as a proxy for new housing supply) and GDP (see Box A for details). The key idea behind our methodology is to identify demand shocks that push prices and quantities in the same direction, which allows us to trace the slope of the supply curve. Housing supply elasticity is then calculated as the ratio of the percentage change in building permits to the percentage change in house prices over the first quarter following a demand shock.[4]

Our elasticity estimates specifically measure the supply elasticity of new housing stock. As building permits capture the flow of the housing stock, our supply elasticity estimates implicitly assume that the marginal new addition to the housing stock determines the price of housing.[5] The reliance on flow measures to compute supply elasticities is common in the literature. For example, Mayer and Somerville (2000b) use housing starts, while Caldera and Johansson (2013) use residential investment.

3 See Annex A, which highlights the importance of rents in CPI baskets. In many economies, rents are also used to measure the costs of owner-occupied housing.

4 We also experimented with elasticity measures based on the cumulative responses over four quarters following a demand shock. We found that the properties of these elasticities do not differ much from those based on the contemporaneous responses.

5 The flow of new housing can even capture changes in the per capita housing stock, for example due to population growth or the need to replenish depreciated housing units (Foote et al (2021) and Loewenstein and Willen (2023)).

Box A

## Estimating the time-varying price elasticity of housing supply

*Ryan Banerjee, Denis Gorea, Deniz Igan and Gabor Pinter*①

We construct estimates of the supply elasticity using a time-varying parameter Bayesian vector-autoregression (BVAR) model, similar to Gorea et al (forthcoming). For each country in our sample, we estimate the following regression:

$$y_t = C_t + \sum_{p=1}^{4} B_{p,t} y_{t-p} + \varepsilon_t,$$

where the dependent variable, $y_t$, is a three-by-one vector of variables that includes the quarter-on-quarter growth rates for house prices, building permits② and GDP. We use a four quarter lag structure of the BVAR model to capture dynamics in the three endogenous variables. $C_t$ is a vector of time-varying country-specific intercepts, $B_{p,t}$ is a matrix of time-varying coefficients, and $\varepsilon_t$ is an error term that is normally distributed with a zero mean and a time-varying covariance matrix.

Our estimation sample includes 21 countries (see Graph 2.C for the specific countries in our study). The length of the time series varies across countries. For most advanced economies, the data start in the 1970s or the 1980s, while for emerging market economies, the sample starts in the 1990s and 2000s.

We identify shocks through sign restrictions, adapting the identification restrictions imposed in Baumeister and Peersman (2013) to the context of the housing market. Table A1 summarises the sign restrictions imposed on the impact (within-quarter) responses to shocks of the three endogenous variables used in the model.

### Sign restrictions to identify supply and demand shocks

Table A1

| | House prices | Building permits | GDP |
|---|---|---|---|
| Aggregate demand shock | + | + | + |
| Housing-specific demand shock | + | + | –[1] |
| Housing supply shock | + | – | – |

[1] The results are robust when this is zero instead of negative.

Source: Authors' elaboration.

The identified demand shocks are crucial for our estimates of the price elasticity of housing supply because they enable us to trace out movements along the supply curve. For our main analysis we use aggregate demand shocks, defined as shocks which push house prices, building permits and aggregate demand in the same direction. The sign restrictions additionally identify housing-specific demand shocks which push prices and quantities up, but GDP down. These shocks are assumed to reduce current GDP, as they tend to disproportionately reallocate investment away from other productive sectors of the economy towards residential construction.③ We distinguish between the two types of demand shocks (aggregate and housing-specific) to better capture the effects of monetary policy that would propagate through the economy in the same way as aggregate demand shocks. We further identify housing supply shocks using our model. These are shocks that move house prices in the opposite direction to building permits and GDP.

We use the impulse responses from aggregate demand shocks to compute estimates of the price elasticity of housing supply. As our model features time-varying parameters, we simulate the impulse responses for building permits and house prices to aggregate demand shocks for each country and quarter. We then use the impulse responses at horizon zero to compute the short-run price elasticity of housing supply, defined as the percentage change in building permits divided by the percentage change in house prices. This elasticity serves as our baseline measure of how much new housing supply changes following a change in house prices.④

① The views expressed are those of the authors and do not necessarily reflect the views of the BIS. ② We use building permits to proxy for the quantity of new housing supply because data on permits are available for a longer time period than data on new housing starts or completions for most countries in our sample. For countries where we have reliable data on starts and permits, we find that the two series are especially highly correlated at lag zero or one, suggesting that permits are a reliable predictor of starts and new supply. ③ Estimates of the price elasticity of housing supply derived from housing-specific demand shocks are similar in terms of dynamics to those derived from aggregate demand shocks. In addition, our results are robust to assuming a zero impact on GDP for the housing-specific demand shock. ④ Using the supply elasticities based on housing-specific demand shocks does not materially change our results on the effects of monetary policy on house prices and rents.

An advantage of our methodology is that it can trace the evolution of the supply elasticity over time. This is important because, as we show below, supply elasticities have changed significantly. By contrast, most of the literature has relied on measures that are time-invariant, for example based on geographical constraints (Saiz (2010)) or on regulatory and institutional constraints on land use (Glaeser et al (2005) and Gyourko et al (2008)).[6]

Our estimates indicate that the supply elasticity has declined over the past five decades on average (Graph 2.A). In the 1970s, a 1% increase in house prices would lead to an average increase in building permits of roughly 6%. By the 2000s, the average elasticity across countries had declined to close to 4%. It recovered slightly in the 2010s but remains well below its level in the 1970s. The decline in the elasticity has been attributed to a tightening in residential land use regulation (Gyourko et al (2021)) and to a drop in the productivity of the construction sector (Goolsbee and Syverson (2023)). Both factors can result in new housing becoming more costly and time-consuming to produce.[7]

Our results further indicate that the supply elasticity varies significantly across countries (Graph 2.B). Consistent with previous studies, the average supply elasticity is higher in the United States than in the United Kingdom.[8] Although geographical constraints are often used to capture exogenous differences in supply elasticities (eg Saiz (2010)), the contrast between Switzerland and Austria, two mountainous countries, suggests that geography may not always be a useful proxy. Finally, there are large differences within AEs and EMEs. Among AEs, elasticities are high in Germany and the United States and low in Australia and Ireland. Among EMEs, they are high in Chile and South Korea and low in emerging Europe.[9]

Cross-country differences in house price dynamics provide a useful crosscheck on the validity of our supply elasticity estimates. A priori, one would expect that the more elastic the housing supply, the smaller the change in house prices for a given change in demand. Consistent with this logic, we find that house prices tend to be less volatile in countries with more elastic supply (Graph 2.C, downward sloping orange line). The flat black line, however, shows that the cross-country relationship does not extend to rents, hinting that the dynamics of rents in response to demand shocks, and hence also to monetary policy, might be quite different from that of house prices.

6 A separate strand of the literature estimates the supply elasticity by regressing changes in the housing stock on changes in prices (Mayer and Somerville (2000a) and Green et al (2005)). However, there are endogeneity issues with this older approach because it also captures how exogenous changes in housing supply affect future prices.

7 Aastveit et al (2023) estimate measures of supply elasticity using an instrumental variables regression for 254 US metropolitan statistical areas spanning the housing boom episodes of 1996–2006 and 2012–19. They also find that the elasticity has been declining over time.

8 See eg Malpezzi and Maclennan (2001).

9 Our elasticity estimates for some countries differ from those in Caldera and Johansson (2013) due to differences in sample periods and in how new housing supply is measured, ie building permits vs residential investment. See Duca et al (2021) for a more detailed discussion of approaches used in the literature to calculate the supply elasticity and the reasons behind differences in elasticity estimates.

Housing supply elasticity has fallen, but there are large cross-country differences[1] Graph 2

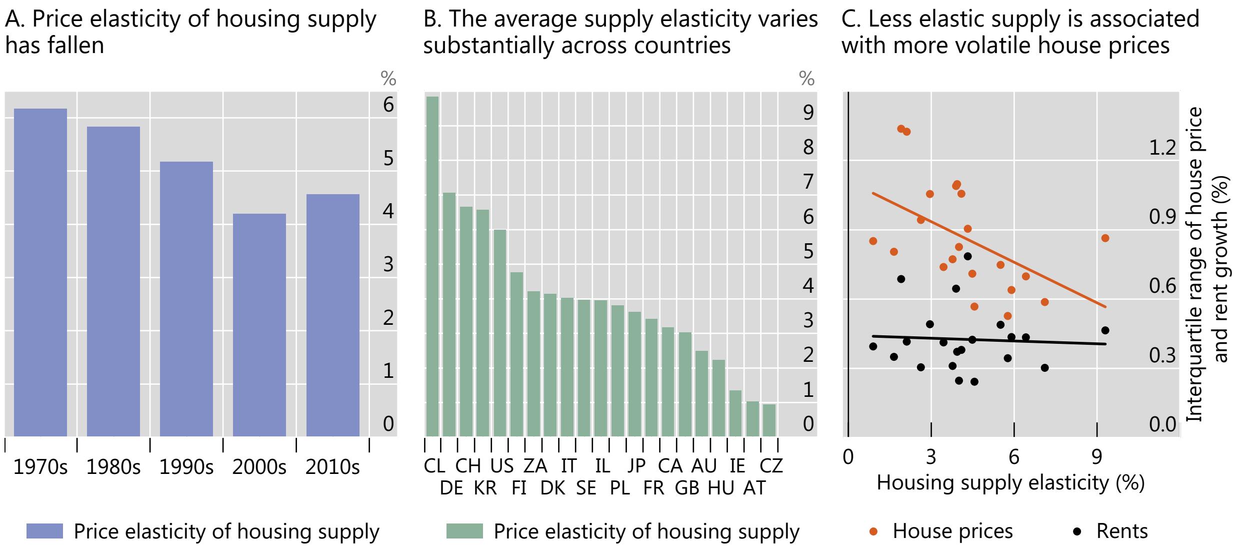


[1] See technical annex for details.

Sources: OECD; LSEG Datastream; Macrobond; national data; authors' calculations.

## Monetary policy transmission to house prices and rents

### The influence of housing supply: evidence from a panel of countries

To estimate how the supply elasticity affects the transmission of monetary policy to house prices and rents, we adapt the local projections (LP) methodology of Jordà (2005) to a panel of countries. Our sample includes 21 AEs and EMEs, covering the period between 1980 and 2019. Specifically, we estimate the following regressions for country $c$ at month $t$ up to horizon $h = 0, \dots, 24$:

$$y_{c,t+h} - y_{c,t-1} = \alpha_c^h + \boldsymbol{\beta_1^h} \cdot \boldsymbol{I_{c,t}} \cdot MP_{c,t} + \sum_{l=1}^{12} \beta_2^h X_{c,t-l} + \varepsilon_{c,t+h},$$

where $y_{c,t+h} - y_{c,t-1}$ is either the cumulative change in the logarithm of the house-price-to-rent ratio, the cumulative percentage change in house prices or the cumulative percentage change in rents; $\alpha_c^h$ is a country fixed effect which controls for time-invariant country characteristics such as land availability. $X_{c,t-l}$ is a set of country-specific controls dated $t - 1$ and earlier which includes changes in the dependent variable, GDP growth, the change in the policy rate, CPI inflation excluding housing, population growth and population density. When estimating the effect of monetary policy on house prices, we also include past changes in rents in the control set. Similarly for the effect on rents, we include past changes in house prices in the control set.[10]

To capture the impact of monetary policy on house prices and rents, we isolate changes in policy rates that are not correlated with the state of the economy.

[10] We do not control for past changes in house prices or rents when estimating the effects on the price-to-rent ratio, as prices and rents are highly correlated with past changes in the dependent variable, which is included as a control in $X_{c,t-l}$.

Specifically, we use monetary policy surprises, $MP_{c,t}$, from Choi et al (2024). These measures of interest rate surprises ensure that our estimates capture the impact of monetary policy itself, rather than the influence of other factors such as the state of the economy. The objects of interest are the coefficients $\boldsymbol{\beta_1^h}$, which vary across low and high groups for the price elasticity of housing supply. Specifically, we construct variable $\boldsymbol{I_{c,t}}$, which is a vector of indicator variables that sort the countries in our panel regression based on whether the country-specific mean elasticity in each decade is above or below the median in the cross section. The regression coefficients are plotted as impulse response functions with confidence bands based on standard errors clustered by country.

We find that the monetary policy impact on the house-price-to-rent ratio depends on the supply elasticity. Our estimates suggest that when supply is less elastic, the price-to-rent ratio increases by approximately 5% two years after a surprise 100 basis point fall in interest rates (Graph 3.A). When supply is elastic, the same monetary policy easing has a weaker effect, raising the price-to-rent ratio by around 2% in the first year, with the effect becoming statistically insignificant thereafter.[11]

This difference in the impact of monetary policy is almost entirely accounted for by its larger impact on house prices. In countries where the housing supply is elastic, a 100 basis point reduction in policy rates results in only a modest 1.5% increase in house prices, which becomes statistically insignificant within a year (Graph 3.B). By contrast, when housing supply is less elastic, the increase in house prices is much stronger and long-lasting, with house prices increasing by around 5% two years after the rate cut (Graph 3.C).[12] We also found that a tightening of monetary policy does not exhibit any clear asymmetry compared with a loosening in low vs high supply elasticity countries. House prices react by more when supply is inelastic both when interest rates rise and when they fall.

Monetary policy has a much weaker impact on rents, regardless of the elasticity. Where supply is elastic, point estimates indicate that rents rise by a similar amount to house prices in response to a cut in interest rates, while they decline slightly when supply is inelastic. But in both cases the estimates are largely statistically insignificant.

## The influence of housing supply: potential channels

If rents are sticky, monetary easing can push up house prices relative to rents. Under standard asset pricing logic, the price of a house equals the present discounted value of future rents. If rents do not respond much to monetary easing, a lower discount rate will push up house prices, raising the price-to-rent ratio.

11 Our results are robust to a variety of different specifications, including grouping countries based on their elasticities in a given quarter, rather than over the decade; using growth in housing starts as an alternative measure of housing supply instead of our baseline supply elasticity estimates; and using supply elasticities from housing-specific demand shocks to rank countries into low and high groups.

12 The response peaks at close to 6% three and half years after the monetary policy shock.

Housing supply elasticity affects the transmission of monetary policy to house prices and rents[1]

In per cent

Graph 3

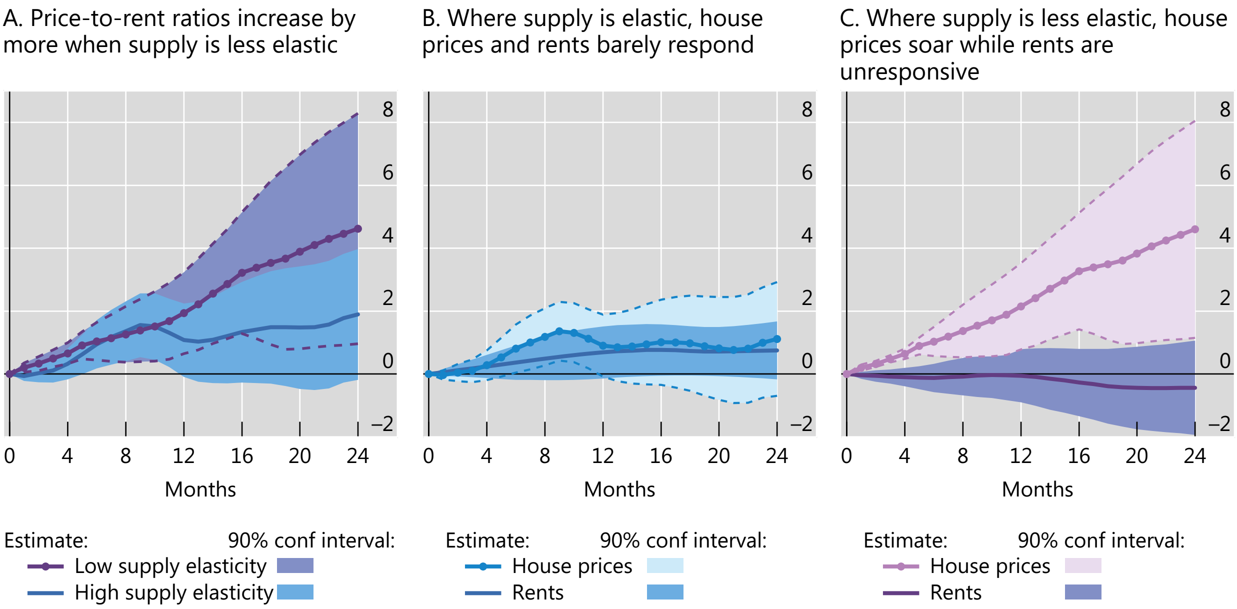


[1] See technical annex for details.

Sources: OECD; Bloomberg; Global Financial Data; LSEG Datastream; Macrobond; national data; authors' calculations.

If supply is inelastic, the overall house price increase following a monetary easing tends to be higher. This is because more of the housing market adjustment to lower interest rates is met by rising house prices rather than by adding new housing units. The fact that house prices adjust by more could additionally result in households extrapolating current house price growth into the future.[13] Strong expectations of future house price growth tend to push up current house prices, which appears consistent with the sustained rise in house prices that we find in the data.

These mechanisms would rationalise the larger rise in house prices, but they struggle to explain why the response of rents is so limited. Assuming substitutability between owning and renting, one would also expect to see higher house prices translating into stronger rent increases when supply is inelastic. This, however, is not what we find in Graph 3. Of course, it is still possible that the adjustment in rents will occur over a much longer horizon than that in house prices.[14] In addition, other forces such as landlords transforming rental units into owner-occupied ones could eventually bring rents more in line with house prices in the long run.

[13] See Adam et al (2024), who find that house price expectations become more extrapolative when the price-to-rent ratio is high.

[14] Another channel could run through credit constraints if they are more binding when supply is less elastic (Albuquerque et al (2024)). Sommer et al (2013) present a model in which the supply and demand in rental markets are jointly determined alongside the demand for owner-occupied housing. When interest rates and down payment requirements fall (eg due to more ample credit supply), the demand for rental units by tenants falls because home ownership becomes cheaper. At the same time, the supply of rental property from landlords increases because investment in rental property becomes more attractive relative to saving in deposits.

## Policy considerations

Our findings have implications for monetary policy in countries with an inelastic housing supply. One implication is that in these countries, the greater impact of monetary policy on house prices may strengthen its transmission through homeowners. The larger changes in house prices can generate especially sizeable changes in the net worth of mortgagors who typically hold significant illiquid assets but have limited liquid wealth. As shown by Cloyne et al (2020), this amplifies their consumption response to monetary policy easing.

That said, when supply is inelastic the strong impact on house prices, but limited traction on rents, could raise questions about the measurement of housing costs in inflation measures targeted by central banks. The public would see large rises in house prices following monetary easing, but at the same time, the component capturing housing costs in CPI would barely move. This is because rents are a major component of CPI, while house prices are typically not included (see Annex A). Communication challenges may be greater in countries that use the rental equivalence approach to impute owner-occupied housing costs in CPI, which effectively places a much larger weight on rents.

More generally, when there is limited monetary policy traction on rents, more of the adjustment burden of bringing inflation to target will fall on other components of CPI. Over the past two years, strong growth in the rental component of CPI has kept inflation high relative to inflation targets in many countries (Banerjee et al (2024)). Policymakers have been confronted by a choice between imposing a fast and costly adjustment on non-housing sectors or tolerating a slower and more gradual return of inflation to target.

Our analysis also highlights the role of broader policies that go beyond the toolkit of central banks. In countries with low or declining supply elasticities, there is scope for supply-side interventions, such as regulatory reforms, that would boost the responsiveness of housing supply to positive aggregate demand shocks. Policy measures that aim to increase the productivity of the construction industry could also help ensure better alignment between supply and demand. Such policies could substantially alleviate the challenges for central banks generated by declining housing supply elasticities.

## References


Aastveit, K, B Albuquerque and A Anundsen (2023): "Changing supply elasticities and regional housing booms", *Journal of Money, Credit and Banking,* vol 55, no 7, pp 1749–83.

Aastveit, K and A Anundsen (2022): "Asymmetric effects of monetary policy in regional housing markets", *American Economic Journal: Macroeconomics*, vol 14, no 2, pp 499–529.

Adam, K, O Pfäuti and T Reinelt (2024): "Subjective house price expectations, falling natural rates, and the optimal inflation target", *Journal of Monetary Economics*, August, 103647.

Albuquerque, B, M Iseringhausen and F Opitz (2024): "The housing supply channel of monetary policy", *ESM Working Papers*, no 59, February.

Andaloussi, M, N Biljanovska, A De Stefani and R Mano (2024): "Feeling the pinch? Tracing the effects of monetary policy through housing markets", IMF *World Economic Outlook*, chapter 2, April, pp 43–63.

Banerjee, R, D Gorea, D Igan and G Pinter (2024): "Housing costs: a final hurdle in the last mile of disinflation?", *BIS Bulletin,* no 89, July.

Baumeister, C and G Peersman (2013): "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market", *Journal of Applied Econometrics,* vol 28, no 7, pp 1087–109.

Bureau of Labour Statistics (BLS) (2024): "Measuring price change in the CPI: rent and rental equivalence", US Bureau of Labour Statistics Factsheet.

Caldera, A and Å Johansson (2013): "The price responsiveness of housing supply in OECD countries", *Journal of Housing Economics*, vol 22, no 3, pp 231–49.

Central Statistics Office (CSO) (2014): *Report of the Group of Technical Advisory Committee on Statistics of Prices and Cost of Living*, Central Statistics Office, India.

Choi, S, T Willems and S Y Yoo (2024): "Revisiting the monetary transmission mechanism through an industry-level differential approach", *Journal of Monetary Economics*, vol 145, July, 103556.

Cloyne, J, C Ferreira and P Surico (2020): "Monetary policy when households have debt: new evidence on the transmission mechanism", *Review of Economic Studies*, vol 87, no 1, pp 102–29.

Dias, D and J Duarte (2019): "Monetary policy, housing rents and inflation dynamics", *Journal of Applied Econometrics*, vol 34, no 5, pp 673–87.

Duca, J, J Muellbauer and A Murphy (2021): "What drives house price cycles? International experience and policy issues", *Journal of Economic Literature,* vol 5, no 3, pp 773–864.

Eiglsperger, M, I Ganoulis, B Goldhammer, O Kouvavas, M Roma and A Vlad (2024): "Owner-occupied housing and inflation measurement", *ECB Statistics Paper Series*, no 47, June.

Foote, C, L Loewenstein and P Willen (2021): "Cross-sectional patterns of mortgage debt during housing boom: evidence and implications", *Review of Economic Studies*, vol 88, no 1, pp 229–59.

Glaeser, E, J Gyourko and R Saks (2005): “Why is Manhattan so expensive? Regulation and the rise in housing prices”, *Journal of Law and Economics*, vol 48, no 2, pp 331–69.

Goolsbee, A and C Syverson (2023): “The strange and awful path of productivity in the US construction sector”, *NBER Working Papers*, no 30845, January.

Gorea, D, B Hofmann, D Igan and J Lee (forthcoming): “Housing supply elasticities across time and space and media sentiment towards new supply”, mimeo.

Green, R, S Malpezzi and S Mayo (2005): "Metropolitan-specific estimates of the price elasticity of supply of housing, and their sources", *American Economic Review,* vol 95, no 2, pp 334–39.

Gyourko, J, J Hartley and J Krimmel (2021): “The local residential land use regulatory environment across US housing markets: evidence from a new Wharton index”, *Journal of Urban Economics,* vol 124, July, 103337.

Gyourko, J, A Saiz and A Summers (2008): “A new measure of the local regulatory environment for housing markets: the Wharton Residential Land Use Regulatory Index”, *Urban studies,* vol 45, no 3, pp 693–729.

Irving Fischer Committee on Central Bank Statistics (IFC) (2006): “Proceedings of the IFC Workshop on ‘CPI measures: central bank views and concerns’”, *IFC Bulletin,* no 24, August.

Jordà, Ò (2005): “Estimation and inference of impulse responses by local projections”, *American Economic Review*, vol 95, no 1, pp 161–82.

Loewenstein, L and P Willen (2023): “House prices and rents in the 21st century”, *NBER Working Papers,* no 31013, March.

Malpezzi, S and D Maclennan (2001): “The long-run price elasticity of supply of new residential construction in the United States and the United Kingdom”, *Journal of Housing Economics*, vol 10, no 3, pp 278–306.

Mayer, C and C T Somerville (2000a): “Land use regulation and new construction”, *Regional Science and Urban Economics*, vol 30, no 6, pp 639–62.

Mayer, C and C T Somerville (2000b): “Residential construction: using the urban growth model to estimate housing supply”, *Journal of Urban Economics,* vol 48, no 1, pp 85–109.

Saiz, A (2010): “The geographic determinants of housing supply”, *Quarterly Journal of Economics*, vol 125, no 3, pp 1253–96.

Sommer, K, P Sullivan and R Verbrugge (2013): “The equilibrium effect of fundamentals on house prices and rents”, *Journal of Monetary Economics*, vol 60, no 7, pp 854–70.

## Annex A: Rents and measurement of housing costs in CPIs

Housing costs are a major part of consumer price indices (CPIs), so the effect of monetary policy on housing markets extends beyond indirect effects on household demand. The overall housing cost component constitutes around 21% of CPI baskets on average (Table A.1). Even after subtracting costs for energy and other utilities, housing cost for rents and owner-occupied housing (OOH) still often make up around 14% of CPI baskets on average.[15]

Rents are the single most important underlying variable feeding into CPI housing costs. Direct expenditures on rents actually paid by households make up around 7% of CPI baskets on average. In addition, rents are often the key variable determining the costs of OOH when it is included in the CPI. In particular, the rental equivalence method is by far the most common way in which statistical agencies include OOH in CPI baskets, with an average weight of around 14% (Table A.1, third column). The rental equivalence method answers the hypothetical question, "How much rent would I demand if I were to rent my home instead of live in it?". To compute the answer to this question, statistical agencies use data on actual rents paid on equivalent housing in the rental sector. Unsurprisingly, there is an almost perfect correlation between actual rents paid and OOH costs as captured by the rental equivalence approach (Graph A.1.A).

In this special feature, we use rent estimates that statistical agencies used to calculate CPI, but there are important measurement issues to consider. CPI rents capture actual rents paid by all households in rental accommodation, which includes ongoing tenancies as well as newly let properties. The CPI rent data, however, have a drawback because sampling methodologies can lead to artificially smooth rents data over short horizons, as it takes time for statistical agencies to capture all changes in rents.[16] To assess the extent of this "fake" persistence in rents, we assess the dynamic relationship between rents on newly let properties and CPI rents for a limited set of countries where we have both series (Graph A.1.B). The co-movement between rents on newly let properties and CPI rents peaks at around six months, suggesting that the CPI rent data will capture the cumulative effect of any rent changes only after about six months. Thus, in our analysis of the impact of monetary policy it is important to look at the impact on rents at horizons longer than six months to assess the overall effects.

While house prices themselves do not directly enter CPI measures of housing costs, they are relatively well correlated with some measures of OOH (Graph A.1.A). The net acquisition approach is the most strongly correlated with house prices. It captures the cost of living in an owner-occupied home by the price of newly built housing excluding the value of the land upon which the house stands. Both the net

[15] Housing costs in the CPI often also include energy and other utility costs, which make up around 7% of the CPI basket on average.

[16] To ensure comparability over time, statistical agencies sample rents from the same dwellings. To limit data collection costs and since rents change infrequently, they do not gather rent data for the same household every month. Instead, they divide the sample into subgroups and typically sample each subgroup every six months. See, for instance, BLS (2024) and CSO (2014) for descriptions of the rent sampling methodology in the United States and India, respectively. This means that if the rent of a dwelling increases one month after the sampling date, it will take statistical agencies another five months (the next sampling date) to incorporate this increase in the index.

acquisition and user cost measures of OOH costs are also strongly correlated with goods price inflation and interest rates, unlike the rental equivalence approach.

Treatment of rented and owner-occupied housing in the target/headline inflation measure in selected countries[1]

Table A.1

| | Housing weight in CPI[2] | *Of which:* | | | |
|---|---|---|---|---|---|
| | | *Rented housing* | *Owner-occupied housing:* | | |
| | | | Rental equivalence | User cost | Net acquisition |
| AR | ✓ (10.5) | ✓ (5.8) | | *Excluded* | |
| AU | ✓ (21.7) | ✓ (6.0) | × | × | ✓ (8.1) |
| BR | ✓ | ✓ | | *Excluded* | |
| CA | ✓ (28.6) | ✓ (7.1) | × | ✓ (18.3) | × |
| CH | ✓ (26.6) | ✓ (15.6) | ✓ (4) | × | × |
| CL | ✓ (7.2) | ✓ (7.2) | | *Excluded* | |
| CO | ✓ (33.1) | ✓ (10.6) | ✓ (14.6) | × | × |
| CZ | ✓ (25.8) | ✓ (3.3) | × | × | ✓ (10.3) |
| DK | ✓ (19.5) | ✓ (9.4) | | *Excluded* | |
| EA | ✓ (14.7) | ✓ (5.6) | | *Excluded* | |
| GB | ✓ (13.1) | ✓ (7.8) | | *Excluded* | |
| HU | ✓(9.1) | ✓(1.5) | | *Excluded* | |
| ID | ✓ | ✓ | | *Excluded* | |
| IL | ✓ (26.8) | ✓ (6.7) | × | ✓ (17.8) | × |
| IN | ✓ (10.1) | ✓ | ✓ | × | × |
| JP | ✓ (21.5) | ✓ (2.5) | ✓ (15.8) | × | × |
| KR | ✓ (17.1) | ✓ (9.9) | | *Excluded* | |
| MX | ✓ | ✓ | ✓ | × | × |
| MY | ✓(23.2) | ✓(17.7) | | *Excluded* | |
| NZ | ✓ (28.0) | ✓ (10.3) | × | × | ✓ (8.7) |
| PH | ✓ (21.4) | ✓(12.8) | | *Excluded* | |
| PL | ✓ (17.4) | ✓ (1.8) | | *Excluded* | |
| SA | ✓(25.5) | ✓(21.0) | | *Excluded* | |
| SE | ✓ (24.3) | ✓ (7.2) | × | ✓ (6.9) | × |
| SG | ✓ (24.8) | ✓ | ✓ | × | × |
| TR | ✓ (14.2) | ✓ (5.1) | | *Excluded* | |
| US (CPI) | ✓ (34.2) | ✓ (10.5) | ✓ (22.5) | × | × |
| US (PCE) | ✓ (18.0) | ✓ (3.6) | ✓ (11.9) | × | × |
| ZA | ✓ (20.8) | ✓ (3.5) | ✓ (13.0) | × | × |

[1] Weights (in per cent) in central bank target/headline inflation measure (based on data availability) in parentheses for 2024 or latest available. [2] The CPI housing component in many countries includes utilities, maintenance and repair, and insurance costs, thus the CPI weight on housing is often greater than the sum of the weights for rented and owner-occupied housing.

Sources: Eiglsperger et al (2024); IFC (2006); OECD; LSEG Datastream; national data; BIS; authors' calculations.

## Rents in CPI: relationship with measures of owner-occupied housing (OOH) and market rents

Graph A.1

A. Correlations of OOH with rents and other prices depend on the methodology[1]

B. A rise in market rents takes time to fully appear in the housing component of CPI[2]

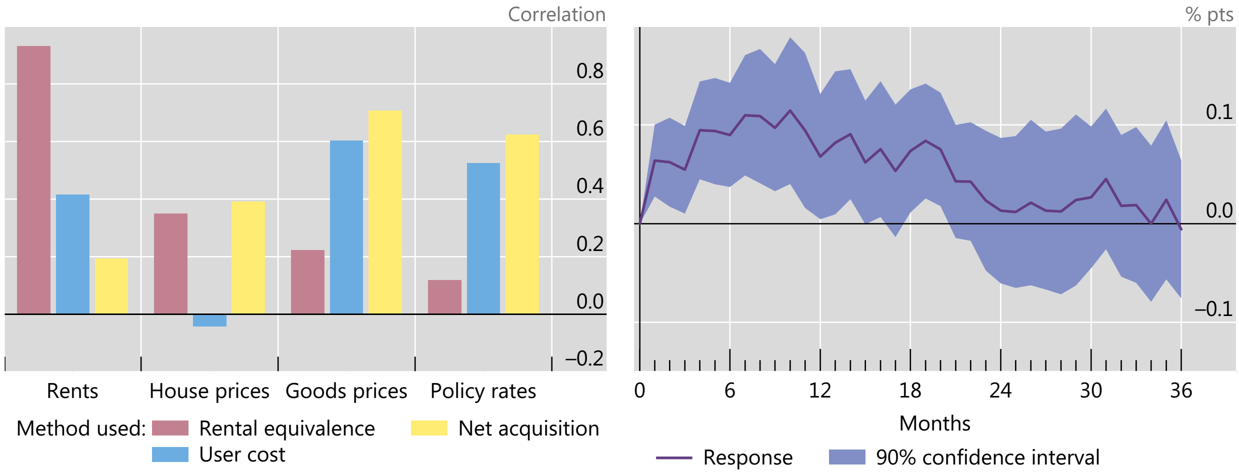


[1] Average correlation coefficient between year-on-year growth rates of the OOH component in CPI and either the year-on-year growth rate in the rental component of CPI (rents), house prices or goods prices or the year-on-year change in policy rates. The sample includes an unbalanced sample of monthly data for 10 economies from January 1961 to May 2024. [2] Impact of a 1% rise in market rents on the housing component in CPI (H-CPI) growth based on local projections that control for lagged growth of H-CPI over the previous month and 12 months and for lags of market rent growth. Based on standardised measures of the housing, services excluding housing, and goods components of CPI, thus can differ from national measures due to the exclusion of energy and other utilities from the housing component. The sample includes an unbalanced panel of monthly data for 12 economies from March 2000 to April 2024.

Sources: OECD; LSEG Datastream; Macrobond; national data; authors' calculations.

## Technical annex

Graph 1.A: GDP-PPP-weighted average across 11 advanced economies (AEs) and six emerging market economies (EMEs).

Graph 1.B: GDP-PPP-weighted average across 14 AEs and seven EMEs. The house-price-to-rent ratio is computed for each economy based on house price indices that use prices for both new and existing dwellings and rent indices constructed using CPI data.

Graph 2.A: GDP-PPP-weighted average across 21 countries (shown in Graph 2.B), computed using a smaller set of countries when data are not available.

Graph 2.B: The average is computed based on varying time periods for each country, depending on data availability. As a consequence, countries for which the data start later in the sample period 1970–2019 (eg Austria and Czechia) tend to have lower elasticities, in line with the evidence shown in Graph 2.A.

Graph 2.C: The sample includes the countries shown in Graph 2.B. For any given country, each dot shows the average estimated supply elasticity and the interquartile range of monthly growth of house prices or rents, deflated by CPI excluding housing. The orange fitted line indicating the relationship between supply elasticity and house price volatility is significant at the 5% level.

Graph 3: Estimates are derived from local projections that regress the cumulative percentage change in the logarithm of the house-price-to-rent ratio, house prices or rents on monetary policy surprises of 100 basis points. The sample includes an unbalanced panel of 21 economies from 1980 to 2019. Economies are categorised as having low (high) supply elasticity if the decadal elasticity of housing supply is lower (higher) than the cross-sectional median.

Bryan Hardy
*bryan.hardy@bis.org*

Patrick McGuire
*patrick.mcguire@bis.org*

Goetz von Peter
*goetz.von.peter@bis.org*

# International finance through the lens of BIS statistics: the geography of banks' operations[1]

*Global banks differ substantially in how they structure their operations across host countries. This feature draws on the BIS international banking statistics to distinguish between "international" bank structures, where positions with counterparties abroad are mainly booked from banks' home offices, and "multinational" structures, where they are booked in affiliates in host countries. It then assesses how differences in these structures shape banks' funding patterns, and what they imply for regulatory oversight and, ultimately, financial stability.*

*JEL classification: F34, F36, G21, F23, F31, F36, G15.*

Banks' global operations are shaped by their history and the economic and regulatory environments in which they operate. Japanese banks rode the post-war industrial boom in Japan to become dominant lenders in cross-border banking by the late 1980s. In the 1990s, Spanish banks expanded their global footprint, mainly in Latin America, with a model that prioritised stand-alone affiliates. For their part, other European banks and US banks embarked on a merger spree throughout the 1990s that transformed them into universal banks with offices in major financial centres that serve all types of clients in every corner of the world. The imprints of these histories are still discernible in banks' global operations today.

Operating abroad exposes banks to complexities and risks that do not arise in purely domestic banking. Banks that lend exclusively at home tend to deal mainly in the domestic currency, face a single regulatory authority, enjoy access to liquidity from the home central bank and compete with other banks in a familiar financial landscape. By contrast, internationally active banks compete on a global stage, face multiple regulatory authorities and deal in various currencies where access to liquidity is not always guaranteed. Banks' *geography* – ie their organisational structure across countries – has bearing on how resilient they are to economic shocks and is thus of first-order importance for financial stability.

This feature is a primer on how the BIS international banking statistics (IBS) can be used to gain insights into the geography of banks' global operations. These statistics track the outstanding assets and liabilities of banks of a given nationality in

[1] The views expressed are those of the authors and do not necessarily reflect those of the Bank for International Settlements. This is the fourth feature in a series showcasing the BIS international banking and financial statistics, following McGuire et al (2024a, 2024b) and Hardy et al (2024). We thank Iñaki Aldasoro, Claudio Borio, Torsten Ehlers, Gaston Gelos, Robert McCauley, Benoît Mojon, Andreas Schrimpf, Hyun Song Shin and Sonya Zhu for their helpful comments, and Swapan-Kumar Pradhan for excellent research assistance.

*Key takeaways*

- *The BIS international banking statistics reveal banks' structures across home and host country locations, thus providing a view of banks' consolidated operations with geographic detail.*
- *Banks with an international structure mainly lend cross-border from their home country. Multinational banks use affiliates in host countries and allocate funds across the group via intragroup positions.*
- *Bank structure shapes the profile of risks that banks face as well as their ability to withstand funding shocks, access foreign currency liquidity and comply with home and host country regulation.*

various locations where they have offices, eg the balance sheet of the offices in Japan of banks headquartered in the United Kingdom ("UK banks"). This helps to shed light, for example, on how intragroup positions (ie "internal markets") enable funds raised in an office in one location to be used by offices elsewhere.

Banks headquartered in 15 countries, which account for the bulk of cross-border banking activity, operate with somewhat distinct organisational structures. Banks with an *international* structure borrow and lend abroad mainly from their home office (eg Japanese banks), drawing on their domestic deposit base. Others adopt a *multinational* structure, operating branches and subsidiaries in multiple jurisdictions. The multinational structure, in turn, can take different forms: *centralised* banks pool funds at major offices (eg in financial centres) and deploy them across the banking group (eg Swiss banks), whereas *decentralised* banks have autonomous subsidiaries that raise deposit funding to finance local assets (eg Spanish banks) (CGFS (2010)).

Each structure is associated with trade-offs and risks. The pooling of resources in a global treasury in international or centralised multinational structures tends to be cost-efficient and robust to local shocks in host countries but can make banks more vulnerable to disruptions in wholesale funding markets. By contrast, a decentralised multinational bank that raises deposits autonomously in each host location where it lends forgoes the benefits of centralised pooling but may be insulated from disruptions in global funding markets.

Episodes of financial volatility and elevated country risk have exposed the importance of bank structure for financial stability. For example, during the Great Financial Crisis (GFC) of 2008–09, centralised multinational banks that had relied extensively on wholesale funding markets were particularly vulnerable when those markets seized up. Banks with autonomous subsidiaries were more insulated from those disruptions. With post-crisis deleveraging and regulatory reforms, banks gravitated towards international or multinational decentralised structures.

Banks' intragroup cross-border funding has played a supportive role during such stress events, but it can also be a vulnerability. For example, when non-US banks faced difficulty in rolling over dollar funding during the GFC and in March 2020, extraordinary policy measures in the form of central bank swap lines were required to restore market functioning. Those banks with affiliates in countries that had access to dollar liquidity were able to deploy funds to affiliates elsewhere, leading to a surge in intragroup cross-border dollar flows. At the same time, multinational banks that rely extensively on cross-border intragroup funding are vulnerable to changes in regulation or macroeconomic downturns in funding locations that may impede the release of funds for use elsewhere.

This feature uses the IBS to first characterise banks' global operations, distinguishing between those with international and multinational (de)centralised structures. It then examines the importance of intragroup funding and how banks' structures affect their resilience to global funding shocks. The final section concludes.

## The architecture of global banking

Historically, many banks that engaged in transactions with counterparties abroad often did not establish a foreign presence. In the 19th century, merchant or investment banks, eg those started by the Morgans, Rothschilds and Barings, were engaged in trade finance and foreign bond issuance; they undertook this activity without establishing affiliates abroad (Jones (1992)).[2] At the same time, British overseas banks established branches in the colonies. A second wave of multinational banking was led by US banks in the 1960–70s, to serve corporates, exploit regulatory differences and run wholesale market operations in financial centres.

The organisational structures seen today have been shaped by domestic economic booms, recurrent financial crises and subsequent regulatory changes, and waves of bank mergers and acquisitions. Riding the post-war economic boom in Japan, low capital costs enabled Japanese banks to lever up and expand internationally. They facilitated the export of capital from Japan to borrowers elsewhere, driving a surge in interbank lending in the second half of the 1980s. Concerns that Japanese banks were undercapitalised during this period were an impetus for the first Basel capital accord in 1988 (Ito and Hoshi (2020)). For their part, several European and US banks that operate today emerged as global giants from the merger and acquisition wave that began in the late 1980s and culminated in the GFC.[3] These banks established extensive international networks with major operations in key financial centres such as London and New York.

### The structure of banks' consolidated global operations

The IBS are a unique source of information for understanding banks' organisational structures across countries. The IBS are aggregate data that group banks by residence (where they operate) and by nationality (where they are headquartered). For each of these perspectives, banks' assets and liabilities are broken down by the sector and country of residence of its counterparties (see Box A). The IBS can thus be used to deconstruct, at the level of national banking systems, banks' consolidated balance sheets according to the geography of their operations across countries. This helps to characterise the structure of banks' operations, as discussed below.

[2] Banks can also rely on correspondent banks to transact with counterparties abroad.

[3] To name but a few, Deutsche Bank acquired Morgan Grenfell Group in the United Kingdom in 1989; Bankers Trust in the United States in 1999; and Scudder Investments, a US asset management firm, in 2002. UBS/SBC acquired SG Warburg plc in London in 1995. In 1997, it acquired Dillon, Read & Co, an investment bank in New York, and later merged with PaineWebber (in 2001). Credit Suisse increased its holdings in First Boston in 1990 and then reorganised into CSFB in 1996–97. Barclays created an investment banking operation in 1986, which subsequently developed into Barclays Capital. In 1995 Barclays purchased the fund manager Wells Fargo Nikko Investment Advisors, which was integrated with BZW Investment Management to form Barclays Global Investors.

Box A

## Geographic dimensions in BIS international banking statistics

The BIS international banking statistics (IBS) provide several complementary perspectives on international banking activity. This box puts the locational (LBS) and consolidated (CBS) versions of the IBS in their geographical context, using the country dimensions introduced in McGuire et al (2024a).

Consider a bank (eg lender *l*) that transacts with a borrower in country *c* (Graph A1). The arrow (*l* → *c*) represents positions between them based on the *residence* view, ie the countries in which the lender and borrower reside. By contrast, viewing banks from a nationality perspective (*n* → *c*) yields a consolidated view of the lenders' exposures to a particular country (Hardy et al (2024)). The nationality view groups the balance sheets of resident units with those of the non-resident affiliates they control (with intragroup positions netted out).

### Country dimensions in locational and consolidated statistics[1] — Graph A1

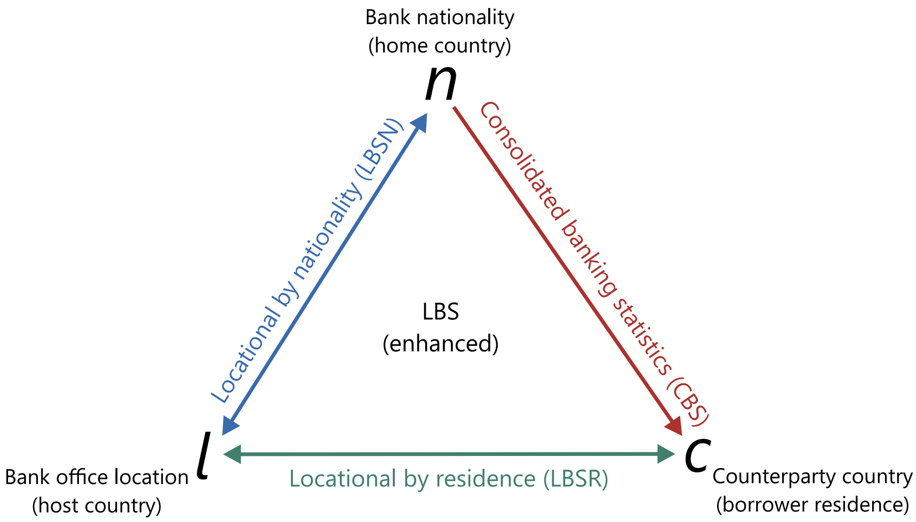


[1] The nodes (*n, l, c*) represent the three geographical dimensions of bank and counterparty available in the BIS international banking statistics (IBS). The arrows in colour depict how the country dimensions relate and which IBS data set covers this perspective; eg the Locational banking statistics by nationality (LBSN) record banks' international claims and liabilities booked in location *l*, broken down by bank nationality *n*.

Source: Authors' elaboration.

The IBS capture all three dimensions at the country level. To illustrate, suppose a German bank books cross-border claims in its London office on counterparties in Brazil. How are these claims reflected in IBS data sets?

- The CBS track banks' worldwide consolidated claims on counterparties in a particular country. The CBS thus capture the *bank nationality* and *counterparty country* dimensions (**red** in Graph A1). The claims of the German bank in London on borrowers in Brazil are aggregated with German banks' total consolidated claims on Brazil, regardless of where the claims are booked (no bank location dimension).
- The LBS by residence (LBSR) track the claims of all bank affiliates located in one jurisdiction on counterparties in each country; the LBS capture the *bank location* and *counterparty country* dimensions (**green**). The German bank's claims on Brazil booked in London appear together with those of all banks in the United Kingdom, regardless of where they are headquartered (no bank nationality dimension).
- The LBS by nationality (LBSN) in every country record bank balance sheets grouped by their home country; they capture the *bank location* and *nationality* dimensions (**blue**). The German bank's claims booked in London on borrowers in Brazil are aggregated with all cross-border claims that German affiliates book in the United Kingdom, regardless of where the borrower abroad resides (no counterparty country dimension).
- The enhanced LBS available since 2013 combine LBSR and LBSN to capture all three dimensions simultaneously: *bank nationality*, *bank location* and *counterparty country*.

A banking group may have offices in multiple host countries that serve different functions. Some offices are sources of funding for the consolidated bank group, to be on-lent elsewhere. By contrast, other offices are *local intermediaries* that fund and lend *within* a host country, or *international intermediaries* that engage mainly in cross-border funding and lending (Box B). In some host jurisdictions, the regulatory and/or economic environment favours one office type over others. A banks' overall consolidated structure will reflect the mix of office types that make up the whole.

A useful starting point for characterising banks' operations is to examine this mix, focusing on *where* banks book their foreign claims and liabilities (Graph 1).[4] At one end of the spectrum are Japanese and German banks, which have large home offices that lend to the rest of the world. Most of their foreign claims are booked as cross-border positions by their home offices in Tokyo and Frankfurt (left bars, blue area). Japanese banks fund around half of their foreign claims via net borrowing from domestic depositors (right bars, purple area).

At the other end of the spectrum are banks that engage mainly in local intermediation in multiple host countries. Spanish banks stand out as having the largest share of their foreign positions booked as local positions in the country of the borrower (Graph 1, green bars). This reflects a preference of banks and supervisors at home and in host countries (eg Latin America) to fund offices in those countries through local deposits (CGFS (2010); Argimón (2019)). Australian banks, with their large presence in New Zealand, have a similar structure.

The geographic patterns suggest that banking systems can be broadly characterised along two dimensions. The first captures banks' local presence across countries (Graph 2.A). Banks with an ***international structure*** run business mainly out

Geography of banks' gross foreign claims and liabilities[1]

End-June 2024 Graph 1

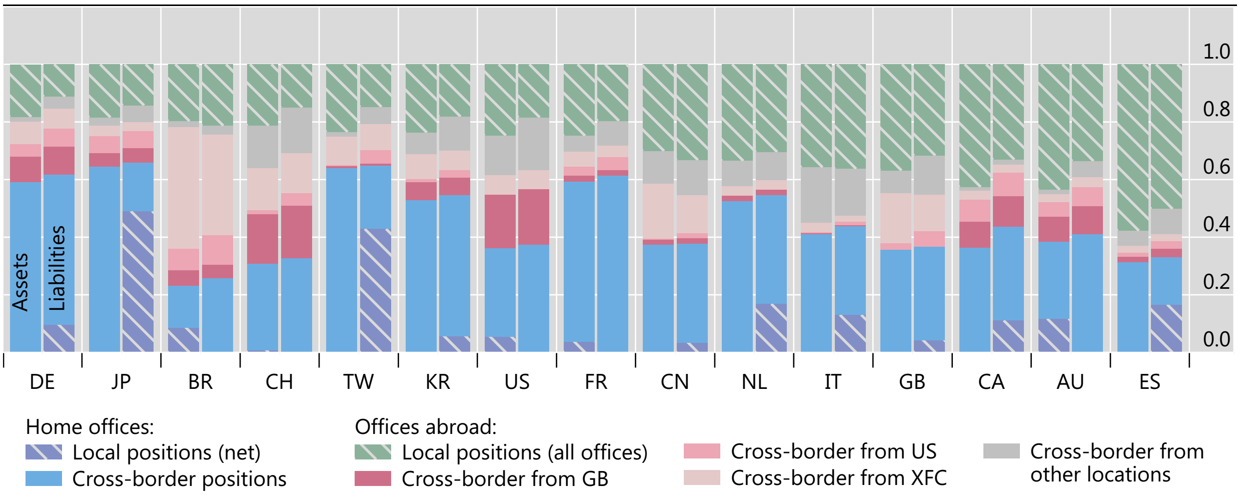


[1] By bank nationality (x-axis), booking location and type. The panel divides total foreign positions into local (hatched bars) and cross-border (solid bars) positions for banks headquartered in the countries listed on the x-axis. For each bank nationality, left (right) bars show assets (liabilities). Cross-border positions are further broken down by office location. The size of each bar indicates the share in total foreign positions. "Local positions" = positions (in all currencies) vis-à-vis residents of the host country; for home offices, only the net position (assets minus liabilities). XFC = cross-border financial centres. See technical annex for details.

Sources: US call reports; BIS consolidated banking statistics; BIS locational banking statistics by nationality; authors' calculations.

[4] The analysis defines home/domestic and foreign with respect to a bank's country of headquarters; the euro area is treated as 20 distinct countries. See also the glossary of terms in Annex A.

of their home country (eg Japanese, German and French banks). Accordingly, they book a low share of positions abroad (dots to the left) and concentrate their positions on a few booking locations (dots towards the top). Other banks have a ***multinational structure***, with sizeable foreign branches and subsidiaries in multiple jurisdictions. Spanish, Swiss, UK and US banks thus appear in the opposite quadrant (dots to the bottom right), since they raise most funds abroad in multiple host countries.

While the international structure is necessarily centralised, the multinational structure can also be decentralised. This second dimension is about the internal organisation of a banking group (Graph 2.B). ***Centralised*** banks pool funds at headquarters or at offices in financial centres and deploy them elsewhere via intragroup positions. Swiss, German, US and Brazilian banks stand out in this regard. ***Decentralised*** banks let their affiliates raise funds autonomously to finance assets in each country (CGFS (2010)). This often involves a high share of local intermediation through retail subsidiaries.[5] Accordingly, Spanish, Italian and Australian banks rely more on local intermediation than on intragroup funding (bottom right).

A bank's structure is rooted in its broader business model – subject to regulatory constraints. Banks with an international structure often rely on a large domestic deposit base to shift capital out of their home countries (eg Japan or Germany), where they centralise control at headquarters. Those with a centralised multinational structure may operate in financial centres outside their home country (eg Swiss banks in wealth management and US banks in investment banking and capital markets). They do so with a central treasury that pools and deploys resources across the group (CGFS (2010); Pokutta and Schmalz (2011); Cetorelli and Goldberg (2012)). More retail-focused decentralised banks with global footprints (eg Spanish banks) often take over local banks and let them operate autonomously (CGFS (2010)).[6]

The structure of banks' global operations, by type[1]

End-June 2024 Graph 2

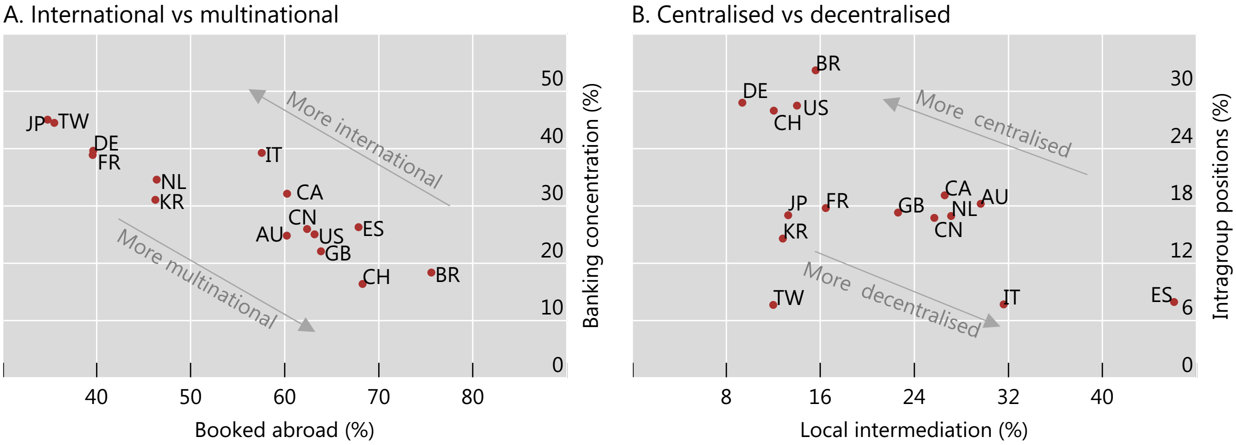


[1] By banking system as shown with ISO country codes. See technical annex for details and Annex A for definitions.

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

[5] Local intermediation gauges the extent to which local claims are funded through local liabilities, rather than through cross-border (eg intragroup) borrowing (see Annex A).

[6] On the pros and cons of a more decentralised structure in which a greater portion of lending to residents of a particular country is funded, managed and supervised by offices in the country, see CGFS (2010). Kamil and Rai (2010) present empirical evidence on the relative stability of banks' local activities in Latin America during the GFC.

Box B

## Office types and the role of banks in facilitating capital flows

Banks' operations across countries can serve a variety of functions, with implications for capital flows. Graph B1 portrays four stylised balance sheet types that depict characteristics captured in the IBS for banks of a particular nationality in a particular host location (eg US banks in Japan). The T-accounts delineate banks' assets and liabilities and the location of the counterparties. Positions with residents of the host country are in grey and cross-border positions are in green for intragroup positions and blue for other counterparties.

Bank offices can be involved in exporting or importing capital, depending on the net position. The home offices of Japanese, German and Swiss banks all **export capital** from the home country and thus have balance sheet structures as in panel A. These offices engage in direct cross-border net lending to counterparties elsewhere and serve as sources of intragroup funding to their offices abroad. Foreign currency assets may be financed with local currency deposits from home country residents, and banks manage the exchange rate risk using foreign exchange swaps. The opposite case is that of a bank office that **imports capital** to the home country (panel B), eg the home offices of Brazilian, US and Australian banks. They can source intragroup funding from their foreign offices or by borrowing from other counterparties (eg from wholesale markets).

Balance sheets by office type[1] Graph B1

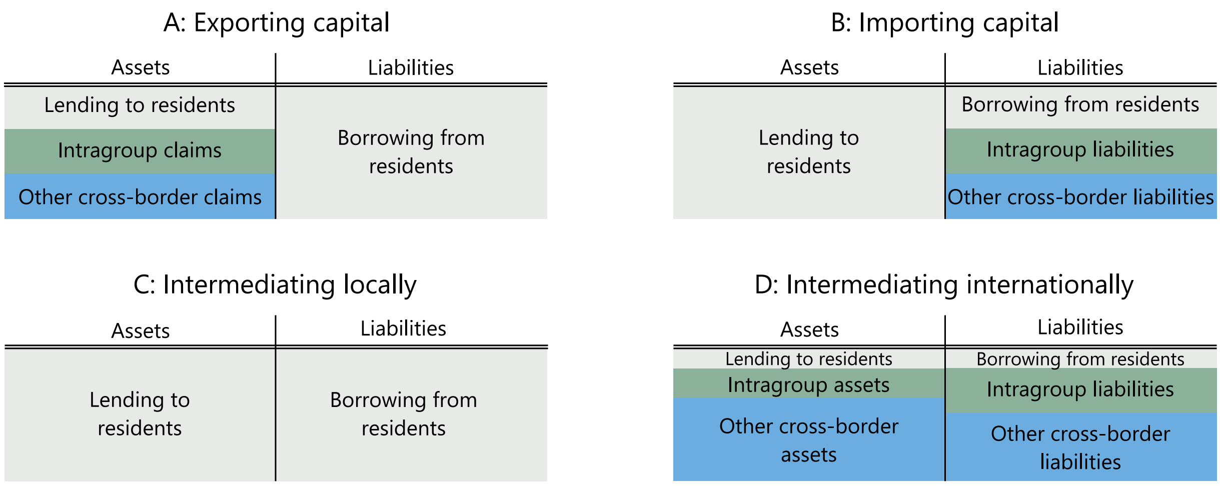


[1] The four stylised balance sheet types shown here can be recreated from IBS data by aggregating bank offices with similar characteristics.

Source: Authors' elaboration.

Offices operating as **local intermediaries** (panel C) typically have local assets denominated in the currency of the host country that are financed with local currency deposits. If under foreign ownership, they are often incorporated as subsidiaries and operate much like domestic banks. They tend to have limited currency mismatches and small cross-border positions and thus contribute little to overall net capital flows to the country.

Finally, there is the **international intermediary** office (panel D), where cross-border positions dominate on both sides of the balance sheet. Such offices serve as routing hubs to shift cross-border funding (eg wholesale funding) to intragroup lending and other cross-border credit to borrowers elsewhere. Examples include banks' offices in major financial hubs (eg London) or in cross-border financial centres (eg the Cayman Islands). Such offices do most business with non-residents and contribute little to domestic credit in the host country, but they do affect the host country's cross-border capital flows and international investment position.

Each structure tends to go with a particular choice between branches and subsidiaries. Unlike subsidiaries, branches do not have their own capital; they operate as an extension of the parent bank's balance sheet. While this makes them easy to manage and integrate in centralised operations, branches may not meet host country regulatory requirements for certain operations, eg taking insured deposits. Decentralised multinational banks therefore rely more on subsidiaries, incorporated as separate legal entities adapted to local regulations and market conditions. Some host countries (eg Brazil, Chile) favour subsidiarisation of foreign banks (Graph 3.A). Banking groups with major operations in such countries (eg Spanish banks) are shown to be decentralised in IBS data. On the other hand, branch activity is more prevalent in some cross-border financial centres (eg the Cayman Islands (KY), Guernsey (GG) and the Bahamas (BS)).

Foreign branches and subsidiaries in BIS reporting countries[1]

In per cent; end-June 2024 Graph 3

A. Share of claims by subsidiaries

B. By sector, currency and position type[3]

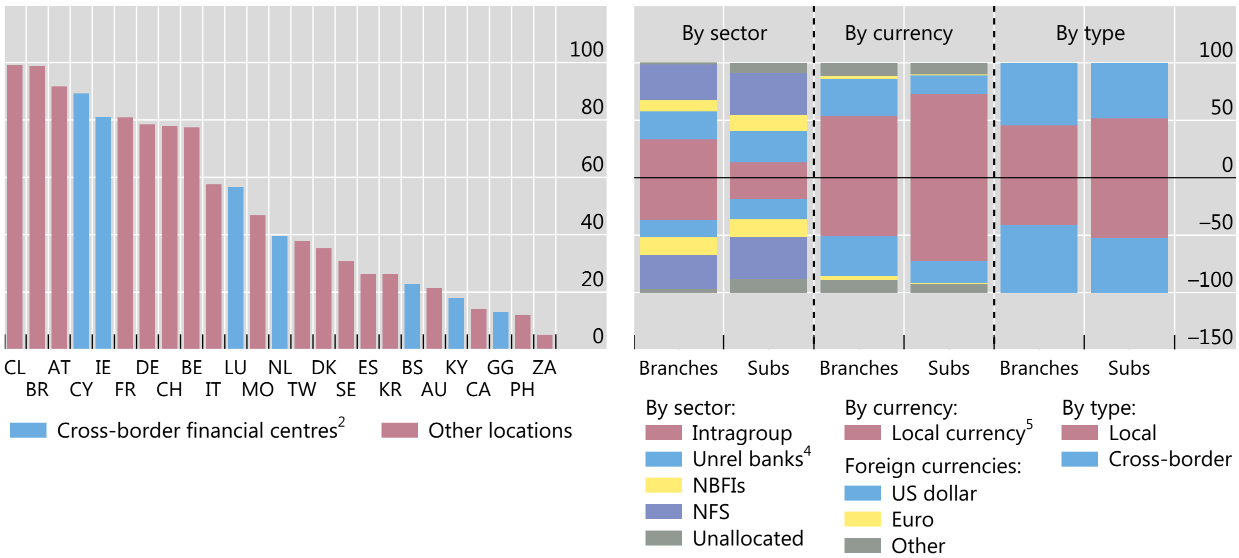


[1] See technical annex for details. [2] Cross-border financial centres as of 2020, defined in Pogliani et al (2022). [3] As a share of total claims (positive bars) and liabilities (negative bars) of foreign branches and subsidiaries in host locations. NBFIs = non-bank financial institutions; NFS = non-financial sectors. [4] Positions with unrelated banks (ie excluding intragroup). [5] Cross-border and local claims positions in the currency of the host country.

Sources: Pogliani et al (2022); BIS locational banking statistics by residence; authors' calculations.

Foreign branches and subsidiaries operate with distinct balance sheet structures. Branches typically have larger interbank positions, particularly intragroup positions, reflecting their integration with the group's balance sheet (Graph 3.B). Subsidiaries, on the other hand, typically have larger positions vis-à-vis the non-financial sector, as being incorporated as subsidiaries facilitates deposit-taking and compliance with local regulatory requirements (Fiechter et al (2011); OECD (2017)). Branches often deal more in foreign currency and more with cross-border counterparties, since they are less often subject to host country limits on foreign exchange and interbank exposures. With this flexibility comes potential instability: compared with foreign subsidiaries, lending by foreign branches tends to be more volatile and more reactive to economic conditions in both their home and host countries (Aldasoro et al (2022)).

Banks' structures have evolved in response to changes in the economic and regulatory environment. Multinational banking was ascendent in the two decades before the GFC, reflected in the rise of local banking. After the emerging market crises of the 1980s and 1990s, banks shifted towards the multinational decentralised structure.[7] As a result, the share of local currency claims in foreign claims on emerging market economies rose from 7% in 1983 to near 40% in the early 2000s (Hardy et al (2024)).[8]

This trend slowed in the 2000s. The introduction of the euro spurred an area-wide interbank market, and European banks ramped up their cross-border investment in asset-backed securities when US markets boomed. This boosted London's position as Europe's financial hub. In emerging markets, cross-border bank flows resumed in the mid-2000s in response to higher yields and US dollar depreciation (Galati et al (2007); CGFS (2009)).

The expansion of cross-border activity from offices in major financial centres proved to be transitory. The GFC exposed weaknesses in banks' overextended balance sheets and forced a broad deleveraging. Disruptions in funding markets disproportionately affected multinational banks that had relied on centralised operations (see the next section). Following the GFC, the international regulatory framework was strengthened (Basel III), and many countries tightened regulation of banks operating in their jurisdiction (Borio et al (2020)). Common measures included requirements to operate as subsidiaries, resolution frameworks and restrictions on business models, including on cross-border positions.[9]

These changes, especially when imposed by host country regulators, discouraged dependence on centralised wholesale funding. Some centralised multinational banks – in particular European banks – reacted by scaling down foreign operations in financial centres, shifting away from a multinational and towards an international structure (McCauley et al (2019)). Other banks gravitated towards the decentralised multinational model (Caparusso et al (2019).[10]

## Implications of bank structure for policy

Bank structure shapes the profile of risks that banks face. For one, it has implications for the resilience of banks' funding models, which depend on their ability to shift funding internally between office locations. For another, the location of banks' affiliates matters in terms of access to central bank liquidity, particularly in foreign currencies. These have bearing on banks' ability to provide credit to the economy and thus are of concern to policymakers.

7 Establishing or acquiring a local bank to borrow and lend locally avoided transfer risk, if not country risk (Hardy et al (2024)).

8 This feature uses "emerging market economies" as a short form for *emerging market and developing economies*, the set of economies not classified as advanced economies in the BIS country grouping convention.

9 See IMF (2015), informed by a survey of bank supervisors in 40 major banking markets, further examined by Ichiue and Lambert (2016).

10 Since 2009, Swiss, German and Dutch banks have downsized their operations abroad, giving their home offices a larger relative role. This shifted them towards the international structure (top left area in Graph 2). By contrast, Korean and Taiwanese banks have shifted in the direction of the decentralised multinational structure (bottom right area).

## Bank resilience during the GFC

International bank credit contracted sharply in the wake of the GFC and again during the European sovereign debt crisis of 2010–12. Many interrelated factors affected the extent of contraction, including banks' business models, their asset quality and funding structures, recessions in countries they operated in, and home and host country regulations.

To disentangle the various factors at play, McGuire and von Peter (2016) examine how major banks' affiliates shed assets across host countries after the GFC. The IBS structure makes it possible to jointly examine the claims and liabilities of banks' foreign offices in different countries, separately from their parent. This in turn enables analysis of the peak-to-trough contraction in assets during the GFC that accounts for factors specific to bank nationality, host country and bank offices (notably their funding mix). Contractions were larger for: (i) offices operating in troubled economies and financial centres (host country effect); (ii) offices that were more exposed to the financial crisis (bank nationality effect); and (iii) offices with a fragile funding mix that relied more on interbank, foreign currency or cross-border funding than on local deposits (funding source effect).[11]

Banks with a decentralised multinational structure fared better in the wake of the GFC. Operating autonomous subsidiaries – ie with local claims funded by local liabilities – can shield banks from shocks to global wholesale funding markets.[12] Banks that had, on the eve of the GFC, booked a higher share of their overall foreign claims locally in local currencies (Graph 4, x-axis) experienced smaller subsequent contractions (y-axis). These affiliates' local operations, which are often retail banking activity, proved to be resilient when global wholesale funding markets seized up.[13] By contrast, affiliates that relied more on other liabilities (eg cross-border, interbank or foreign currency funding) saw larger balance sheet contractions.

11 Hahm et al (2013) and de Haas and van Lelyveld (2014) find similar effects using country-level and bank-level data, respectively. Peek and Rosengren (1997) trace how Japanese banks transmitted financial distress at home to the supply of credit in the United States through their US affiliates. IMF (2015) finds cross-border lending is more volatile and has a stronger transmission of global shocks compared with local lending.

12 At the same time, autonomous affiliates can expose banks to local shocks and limit their role in intermediating capital flows that can benefit host countries.

13 For each country grouping, there is a weak but discernible downward slope in the regression line. That said, there are some clear outliers. For example, Dutch and Belgian banks experienced some of the largest overall contractions, reflecting the breakup of ABN AMRO and Fortis. By contrast, Japanese banks hardly contracted at all, yet they had the lowest ratio of local claims to foreign claims. Japanese banks were less exposed to toxic assets and suffered smaller losses during the crisis than did many European banking systems.

## Local orientation mitigating the contraction in foreign claims around the GFC[1]

Graph 4

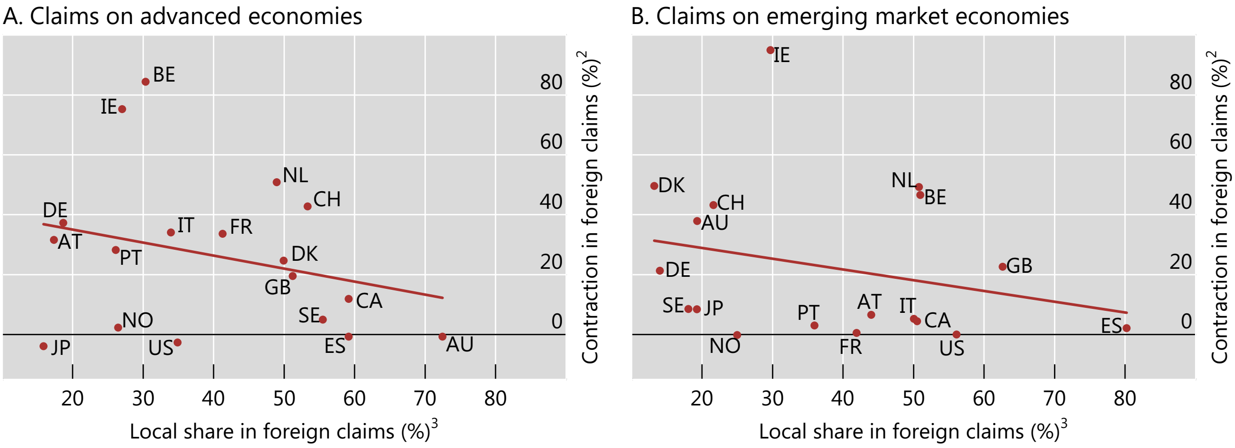


[1] See technical annex for details. [2] Peak-to-trough contractions around the GFC of the foreign claims of the indicated bank nationalities on borrowers in the group of countries indicated in the panel heading. [3] Share of foreign claims booked locally in local currency (by the bank's affiliate there).

Sources: BIS consolidated banking statistics (immediate borrower basis); BIS locational banking statistics by nationality; authors' calculations.

## Internal markets: banks' intragroup funding

Banking groups routinely deploy funds raised in funding markets to investment markets (Cetorelli and Goldberg (2012)). A bank may set up funding operations in countries with a large deposit base or deep bond markets to fund credit by affiliates elsewhere that cannot rely on local deposits. For example, Swedish banks financed their operations in the Baltics in the 2000s through intragroup funding (CGFS (2010)). In the IBS for mid-2024, the United Kingdom and Singapore appear as funding centres – as evidenced by large net intragroup claims (Graph 5.A). By contrast, banks in the United States, Japan and Italy are net users of intragroup funds.

Domestic and foreign banks in a given country use intragroup positions in somewhat distinct ways (Graph 5.B). In France, Japan and Switzerland, for instance, most positions are reported by domestic banks (blue bars), while foreign banks (red bars) play a greater role in countries with financial centres (eg United States, United Kingdom) and in smaller cross-border financial centres (XFC). Domestic banks typically extend net intragroup funding to their affiliates abroad (blue dots), while foreign branches and subsidiaries are net receivers (red dots below zero).

Dollar funding is of particular importance, since the US dollar is the dominant currency across financial markets (McGuire et al (2024b)). Non-US banks hold large asset positions denominated in US dollars and use their network of affiliates around the world to source dollar funding. Less than a quarter (23% at end-2023) of their dollar liabilities are booked by their affiliates in the United States, where dollar funding markets are deep and banks have access to Federal Reserve facilities. The bulk is borrowed either by their home office or by offices elsewhere, in jurisdictions where the dollar is a foreign currency (Aldasoro and Ehlers (2018)).

## Internal markets: banks' cross-border intragroup positions[1]

End-June 2024 Graph 5

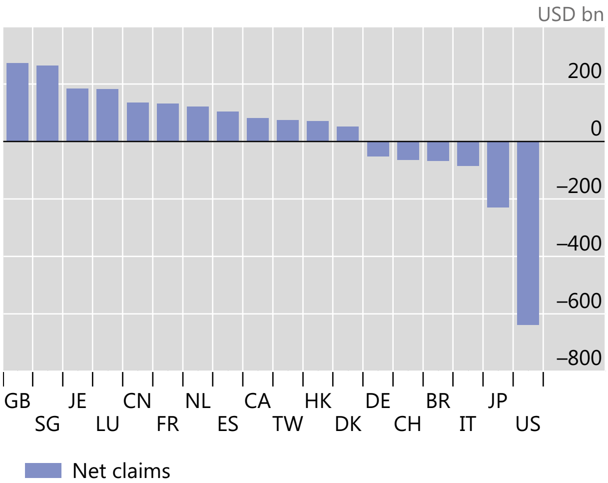


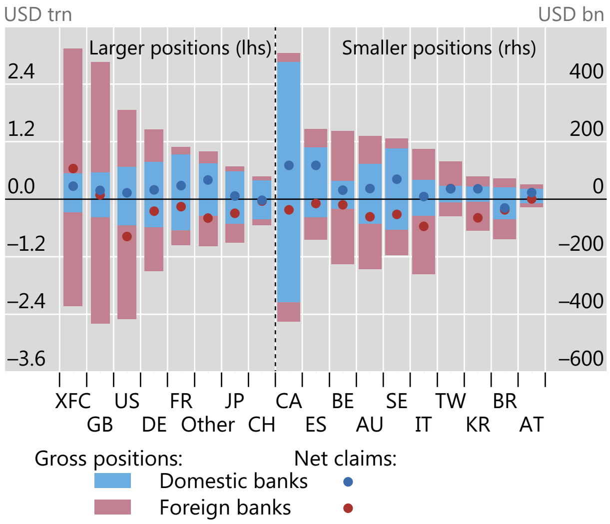


[1] See technical annex for details. Net claims are cross-border intragroup claims minus intragroup liabilities of all reporting banks located in the country shown on the x-axis. [2] Intragroup gross claims (positive) and liabilities (negative) for domestic (blue) and foreign (red) banks. Other represents reporting jurisdictions other than XFC (cross-border financial centres) and those listed individually.

Sources: Pogliani et al (2022); Pradhan and Silva (2019); BIS locational banking statistics (by nationality); authors' calculations.

Intragroup positions are especially important following regulatory changes. Prior to the GFC, non-US banks raised dollar funding from their US operations to lend intragroup to affiliates elsewhere. With US quantitative easing after the GFC and subsequent changes in Federal Deposit Insurance Corporation (FDIC) regulation, intragroup flows reversed as these banks invested heavily in reserve holdings at the Federal Reserve (Kreicher et al (2013)).

Similarly, intragroup positions are essential for reallocating funds during periods of financial stress. During the GFC and again in March 2020, non-US banks tapped dollar facilities at major central banks and used intragroup positions to allocate dollars. Intragroup flows in all currencies explained almost $1 trillion of the outsize expansion in cross-border flows in the first quarter of 2020 (Graph 6.A). The surge in intragroup dollar flows, mostly vis-à-vis affiliates in the United States (Graph 6.B), largely reversed in the second quarter as financial conditions improved (Aldasoro et al (2020)).

These extraordinary movements reflect how funds secured via the Federal Reserve's US dollar swap lines propagated through the global banking network via intragroup positions. By end-March 2020, a total of $358 billion had been drawn from dollar swap lines. In practice, banks draw on the swap line by booking a liability to their home central bank and an intragroup claim on their US affiliate that obtains dollar reserves at the Federal Reserve (Aldasoro et al (2020); McCauley (2024)). Not by coincidence, banks in Japan, the euro area and the United Kingdom, which tapped dollar swap lines the most, reported the largest increases in intragroup dollar claims on their affiliates in the United States (Graph 6.C).[14]

[14] Cetorelli et al (2020b) suggest that the dollar liquidity obtained through the swap lines was an important source of funding for many foreign banks in the United States to handle drawdowns of credit lines by their customers.

Swings in cross-border positions reflect banks' internal markets[1]

In billions of US dollars

Graph 6

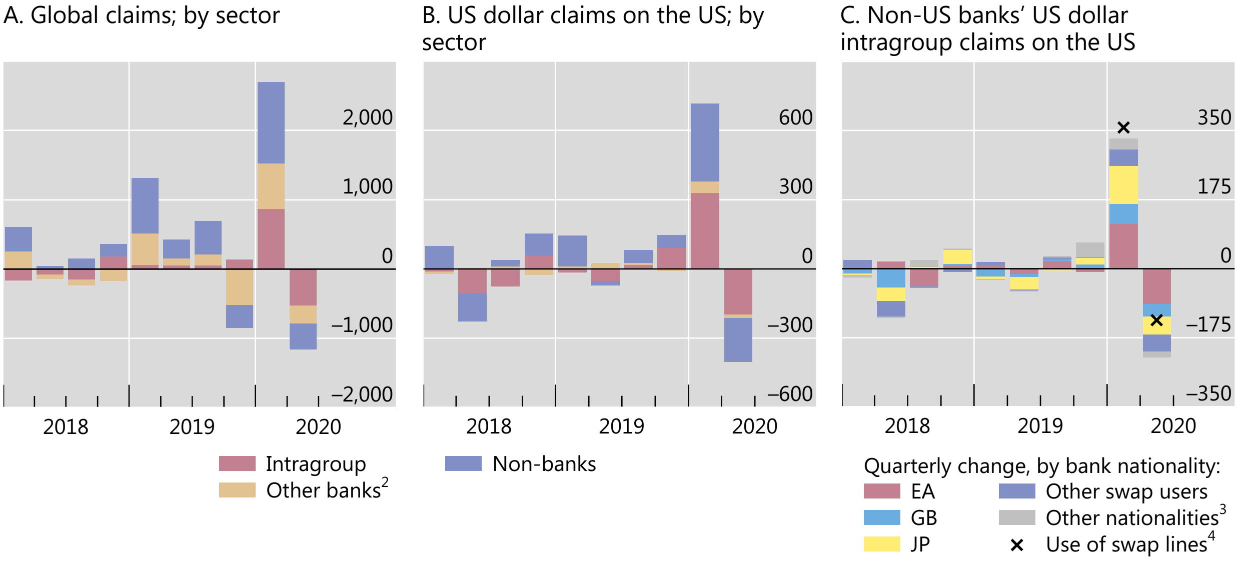


[1] Panel A shows quarterly changes adjusted for breaks in series and exchange rate fluctuations. US dollar claims in the remaining panels require no exchange rate adjustment. See technical annex for details. [2] Includes central banks and banks unallocated by subsector between intragroup and unrelated banks. [3] All other bank nationalities in BIS reporting countries. [4] Use of central bank liquidity swap lines; quarterly change in amount outstanding.

Sources: Federal Reserve Bank of New York; BIS locational banking statistics (by residence and by nationality); authors' calculations.

## Conclusion

Banks' organisational structures reflect their unique histories and business models and have evolved with changes in the regulatory and economic environments in which banks operate. The BIS IBS are a unique source of data on banks' structures since they provide a view of banks' global consolidated operations at the level of the constituent parts. This reveals not only the location of banks' major affiliates but also the location of their funding sources and use of internal markets.

The geography of banks' global operations has important policy implications. It has a bearing on banks' resilience. Supervisory guidance by both home and host regulators can push banks to a decentralised multinational structure, which proved resilient during the GFC. At the same time, a purely local orientation in a host country limits a foreign bank's ability to intermediate capital flows that can benefit host countries. Further, managing access to foreign currency liquidity during stress events is a global challenge that requires cooperation between central banks: banks without affiliates in countries with access to dollar liquidity facilities may face constraints.

## References


Aldasoro, I, C Cabanilla, P Disyatat, T Ehlers, P McGuire and G von Peter (2020): "Central bank swap lines and cross-border bank flows", *BIS Bulletin,* no 34, December.

Aldasoro, I, J Caparusso and Y Chen (2022): "Global banks' local presence: a new lens", *BIS Quarterly Review*, March, pp 31–43.

Aldasoro, I and T Ehlers (2018): "The geography of dollar funding of non-US banks", *BIS Quarterly Review*, December, pp 15–26.

Argimón, I (2019): "Spanish banks' internationalisation strategy: characteristics and comparison", Banco de España, *Economic Bulletin*, no 1, January.

Borio, C, M Farag and N Tarashev (2020): "Post-crisis international financial regulatory reforms: a primer", *BIS Working Papers,* no 859, April.

Caparusso, J, Y Chen, P Dattels, R Goel and P Hiebert (2019): "Post-crisis changes in global bank business models: a new taxonomy", *IMF Working Papers,* no 19/295, December.

Cetorelli, N and L Goldberg (2012): "Liquidity management of US global banks: internal capital markets in the great recession", *Journal of International Economics*, vol 88, no 2, pp 299–311.

Committee on the Global Financial System (CGFS) (2009): "Capital flows and emerging market economies", *CGFS Papers*, no 33, January.

——— (2010): "Funding patterns and liquidity management of internationally active banks", *CGFS Papers*, no 39, May.

de Haas, R and I van Lelyveld (2014): "Multinational banks and the global financial crisis: weathering the perfect storm?", *Journal of Money, Credit and Banking*, vol 46, no s1, pp 333–64.

Fiechter, J, I Ötker-Robe, A Ilyina, M Hsu, A Santos and J Surti (2011): "Subsidiaries or branches: does one size fit all?", *IMF Staff Discussion Notes,* no 11/04, March.

Galati, G, A Heath and P McGuire (2007): "Evidence of carry trade activity", *BIS Quarterly Review*, September, pp 27–41.

Hahm, J-H, H S Shin and K Shin (2013): "Noncore bank liabilities and financial vulnerability", *Journal of Money, Credit and Banking*, vol 45, issue s1, pp 3–36.

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, pp 69–86.

Ichiue, H and F Lambert (2016): "Post-crisis international banking: an analysis with new regulatory survey data", *IMF Working papers,* no16/88, April.

International Monetary Fund (IMF) (2015): "International banking after the crisis: increasingly local and safer?", *Global Financial Stability Report*, chapter 2, April, pp 55–91.

Ito, T and T Hoshi (2020): *The Japanese economy*, second edition, MIT Press.

Kamil, H and K Rai (2010): "The global credit crunch and foreign banks' lending to emerging markets: why did Latin America fare better?", *IMF Working Papers,* no 10/102, April.

Kreicher, L, R McCauley and P McGuire (2013): "The 2011 FDIC assessment on banks' managed liabilities: interest-rate and balance-sheet responses", *BIS Working Papers,* no 413, May.

McCauley, R (2024): "The offshore dollar and US policy", *Federal Reserve Bank of Atlanta Policy Hub,* no 2-2024.

McCauley, R, A Bénétrix, P McGuire and G von Peter (2019): "Financial deglobalisation in banking?", *Journal of International Money and Finance*, vol 94, June, pp 116–31.

McGuire, P and G von Peter (2016): "The resilience of banks' international operations", *BIS Quarterly Review*, March, pp 65–78.

McGuire, P, G von Peter and S Zhu (2024a): "International finance through the lens of BIS statistics: residence vs nationality", *BIS Quarterly Review*, March, pp 73–88.

——— (2024b): "International finance through the lens of BIS statistics: the global reach of currencies", *BIS Quarterly Review*, June, pp 1–16.

Organisation for Economic Co-operation and Development (OECD) (2017): *The conditions for establishment of subsidiaries and branches in the provision of banking services by non-resident institutions*, January.

Peek, J and E Rosengren (1997): "The international transmission of financial shocks: the case of Japan", *American Economic Review*, vol 87, no 4, pp 485–505.

Pogliani, P, G von Peter and P Wooldridge (2022): "The outsize role of cross-border financial centres", *BIS Quarterly Review*, June, pp 1–15.

Pokutta, S and C Schmalz (2011): "Managing liquidity: optimal degree of centralization", *Journal of Banking and Finance*, vol 35, no 3, pp 627–38.

Pradhan, S-K and J-F Silva (2019): "Using mirror data to track international banking", *IFC Working Papers,* no 19, December.

## Annex A: The structure of banks' operations

This annex defines key terms and concepts used throughout the feature, based on the BIS Glossary and the Box in the previous feature in this series (Hardy et al (2024)).

1. **Domestic positions:** claims [15] or liabilities of a domestic bank vis-à-vis residents of the bank's home country.
2. **Cross-border positions:** positions on a non-resident – for example, claim on or liability to a counterparty located in a country other than the country where the banking office that books the position is located.
3. **Local positions:** claims on or liabilities to a counterparty located in the same country as the office that books the positions (the opposite of cross-border positions). Local positions in the home country are domestic positions (see 1.).
4. **International positions:** positions on a non-resident or denominated in a foreign currency. International claims comprise banks' cross-border claims in all currencies plus their local claims denominated in foreign currencies.
5. **Foreign positions:** claims on, or liabilities to, counterparties outside the bank's home country (typically the country of headquarters). Foreign claims comprise positions 2 and 3 above, ie local claims of the bank's offices abroad as well as cross-border claims of the bank's offices worldwide.
6. **Gross foreign positions**: foreign positions plus intragroup positions and net local claims (or liabilities) on home country residents (Graph 1). This, together with gross domestic positions on home country residents, makes up the total of banks' global positions, ie all financial positions on banks' balance sheets.

The various indicators on the structure of banks' operations (Graph 2) can be constructed from foreign positions broken down by booking location. Let $FC_n$ denote gross foreign claims (and $FL_n$ gross foreign liabilities) of banks headquartered in country *n* (nationality), and $FC_{nl}$ denote those booked in country *l* (office location), so $FC_n = \sum_l FC_{nl}$. When foreign claims and liabilities are combined, the share of foreign positions booked in country *l* equals $s_{nl} = (FC_{nl} + FL_{nl})/(FC_n + FL_n)$.

The share ***booked abroad*** is the share of positions booked by offices abroad, ie the share not booked at home, $1 - s_{nn}$. The share booked abroad equals 0 for banks booking all positions in their home country, as would be the case for purely domestic banks or international banks booking cross-border claims from their headquarters.

***Booking concentration*** can be measured by the Herfindahl-Hirschman index, $\sum_l s_{nl}^2$. It equals one if all positions are booked in one location (eg at home) and falls when banks book positions in a greater number of countries in similar amounts.

***Intragroup positions*** are the sum of gross intragroup claims and liabilities across all offices, as a share of gross foreign positions $(FC_n + FL_n)$.

***Local intermediation*** is the minimum of local claims and local liabilities (in all currencies) across all office locations, scaled by overall foreign claims:

$$\sum_l \frac{\min\{LC_{nl}, LL_{nl}\}}{FC_n},$$

where $LC_{nl}$ stands for *local claims* in country *l* booked by banks from *n*, and $LL_{nl}$ likewise stands for local liabilities. This indicator is zero for banks that source liabilities

[15] In the BIS LBS, bank claims comprise (i) loans and deposits, (ii) holdings of debt securities, and (iii) derivatives with a positive market value and other residual instruments (combined). Credit is defined as the sum of (i) and (ii). In the BIS CBS, derivatives are broken out from claims. See Hardy et al (2024).

in one country to fund lending in another. It approaches unity for banks acting as local intermediaries without relying on any cross-border lending or funding.

## Technical annex

All graphs are based on data from the LBS and CBS. While China does not report to the CBS, its data in Graph 1 and Graph 2 are estimated using (incomplete) data available in these two data sets. Banking systems with long historical data for local claims in local currency are BE, CA, CH, DE, ES, FR, GB, IT, JP, NL and US. For the denominator (foreign claims), all reporters are included irrespective of whether a country reported local claims in local currency. The cross-border financial centres (XFC) comprise 23 jurisdictions, of which 14 report locational banking statistics. They are BH, BM, BS, CY, GG, HK, IE, IM, JE, KY, LU, NL, PA and SG. The data for these jurisdictions are shown aggregated together in Graph 1, Graph 3.A and Graph 5.B.

"Emerging market economies" is a short form for *emerging markets and developing economies*, the set of economies not classified as advanced economies in the BIS country grouping convention. Cross-border financial centres are defined in Pogliani et al (2022). The names of jurisdictions corresponding to ISO codes are provided under the abbreviations on pages iv–vii.

Graph 1: All LBS reporting jurisdictions except BH, CN, IN, JE, JP, PA, SG, TR and US report local currency claims and local liabilities of banks' foreign affiliates in their jurisdictions (IN and JP report these local positions only for domestic banks). The consolidated banking statistics on an immediate borrower basis (CBSI) data fill this gap for bank nationalities shown on the x-axis. The CBSI are also sourced for these positions in all countries other than BIS LBS reporting countries. Intragroup claims (liabilities) vis-à-vis non-reporting countries are mirrored as liabilities (claims) of affiliates in these locations. XFC = cross-border positions booked by offices in cross-border financial centres (Pogliani et al (2022)), excluding NL, which is shown separately. Local positions of banks in the United States for Chinese banks are sourced from the US call reports.

Graph 2: Banking concentration, intragroup positions and local intermediation are defined in Annex A. The Herfindahl-Hirschman index is calculated for gross foreign claims and liabilities (both cross-border and local, including intragroup) by individual office location and expressed as a percentage. Booked abroad relates to the percentage of gross foreign claims and liabilities booked by offices outside the home country.

Graph 3: Based on the bank type breakdown (branch vs subsidiary) reported by 35 of 46 reporting countries (does not include BH, CN, GB, JE, JP, MX, NO, PA, SG, TR and US).

Graph 4: For each banking system, peak values are the maximum values observed in the Q1 2006–Q1 2009 period, and trough values are the minimum values in the Q2 2009–Q2 2012 period. Locally booked claims in local currency relate to the average share during the Q1 2006–Q1 2009 period. Taken from McGuire and von Peter (2016).

Graph 5: Domestic banks and foreign banks by host locations. The LBSN is used for the split of domestic and foreign banks. Reporting gaps in bilateral intragroup positions are filled in using mirror data techniques (see Pradhan and Silva (2019)).

Panel A shows net positions of all banks (domestic and foreign banks) in the host locations.

Graph 6: Non-banks in panels A and B include non-banks and those unallocated by sector. Changes in intragroup claims (panel C) are derived from estimated stocks using mirror data reported by two or more jurisdictions for each observation. The approach reconciles banks’ reported claims on related offices in the United States with the intragroup liabilities of those same offices vis-à-vis reporting banks outside the United States. EA banks are those headquartered in any of the 13 BIS reporting euro area countries. Other swap users comprise AU, CH, DK, KR, NO, MX and SG (KR and MX started to use their swap lines in Q2 2020).