

The farsightedness paradox

How central banks need to prepare for an AI future

OLIVER WYMAN
A MARSH BUSINESS

Central banks face a distinctive AI challenge. Their approach to AI adoption must be carefully calibrated, as their institutional legitimacy rests on judgment, accountability, independence, and public trust.

Yet moving too slowly is also risky. The financial system is becoming more complex, more data-intensive, and increasingly shaped by AI-enabled institutions, infrastructures, and threats.

For central banks, the arrival of AI is significant in two fundamental ways. First, it is a tool for the institution: it can enhance macroeconomic analysis, financial supervision, payments oversight, and internal effectiveness. Second, it is a feature of the system they oversee, changing how financial institutions make decisions, how risks emerge and propagate, and how markets and infrastructures behave under stress.

This dual relevance creates a strong case for moving from experimentation to value at scale. Central banks must adopt a deliberate approach to identifying where AI can create institutional value and which capabilities are required to deploy it safely at scale.

The case for central bank action is becoming increasingly compelling. As models evolve rapidly, more powerful AI systems could become a new source of financial-market risk by enabling faster and more autonomous cyber activity. In turn, this could amplify vulnerabilities, intensify market disruption, and accelerate the transmission of stress across interconnected institutions.

AI systems could become a new source of financial-market risk by enabling faster and more autonomous cyber activity.

MYTHOS' IMPACT FROM CENTRAL BANK PERSPECTIVE

The emergence of advanced AI cyber-reasoning systems, such as Anthropic's restricted-access Mythos models, has shifted the AI conversation from productivity to systemic resilience. These models reportedly identify software vulnerabilities and exploit paths at machine scale, dramatically compressing the time between vulnerability discovery and potential exploitation. For central banks, this changes the operational assumptions underpinning cyber resilience, financial stability monitoring, and supervisory expectations.

The relevance to central banking is profound. AI systems capable of autonomous cyber discovery may increase the likelihood of coordinated attacks on payments infrastructure, market utilities, core banking systems, and financial messaging networks. This creates a new category of systemic risk: AI-amplified operational contagion. Historically, cyber risk has been treated as an operational risk issue, but in an AI-enabled environment, it increasingly becomes a macro-financial stability issue.

The implication is that central banks may need to revise supervisory imperatives in three areas:

(1) faster patching and incident response cycles, (2) AI-enabled cyber stress testing across supervised entities, and (3) closer coordination between prudential supervision, cyber authorities, and payments oversight teams.

This also strengthens the case for central banks to develop internal AI capabilities — not only to improve policy effectiveness, but to maintain parity with increasingly AI-enabled financial institutions and threat actors.

© Oliver Wyman

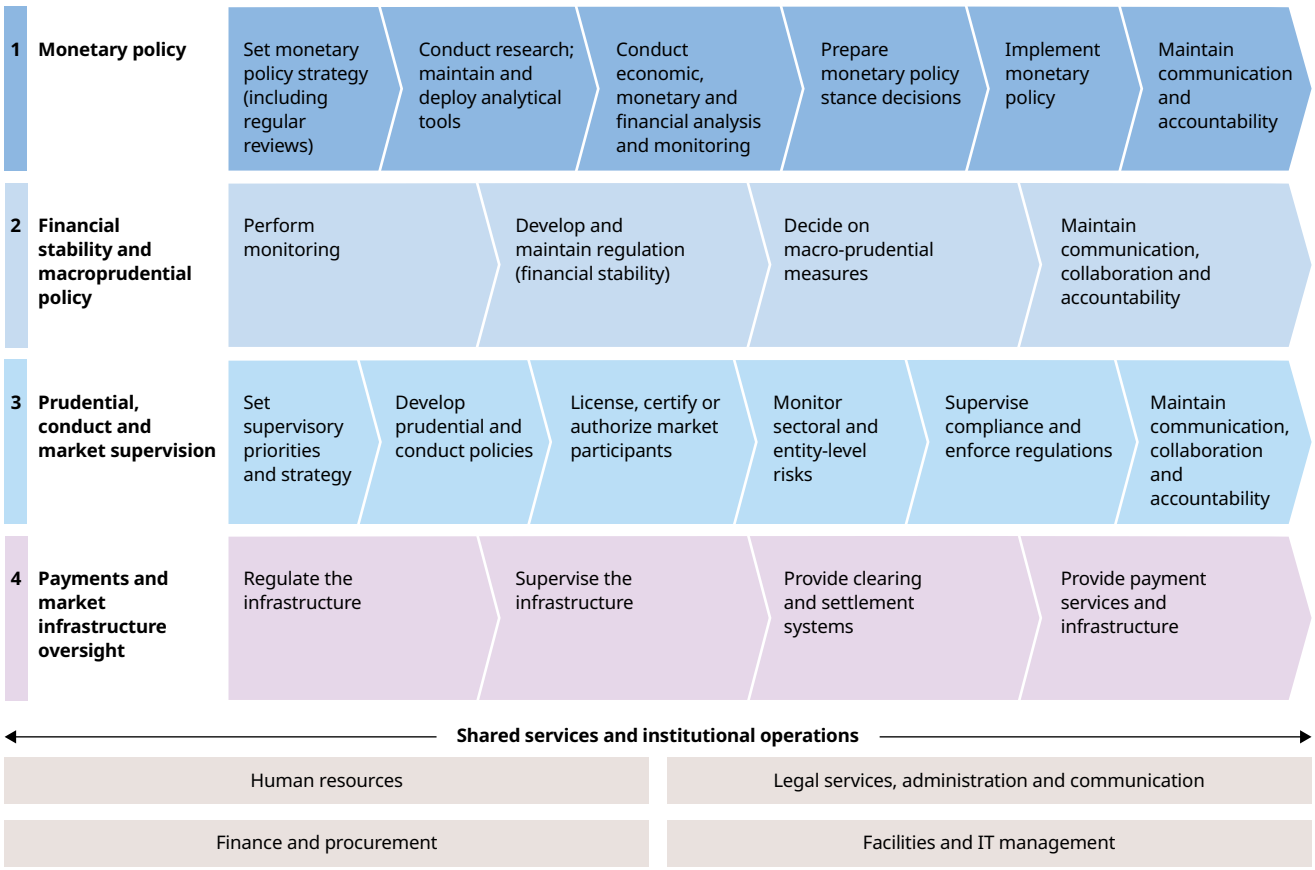
The farsightedness paradox

From a distance, the AI opportunity for central banks looks obvious. Central banks sit on large data assets, conduct complex analysis, monitor vast systems, and make consequential decisions with many moving parts and uncertain elements. That's the perfect environment in which AI thrives.

AI can help with multiple functional tasks: data collection, validation, analysis, decision support, implementation, monitoring, communication, and institutional learning. And it provides speed that humans alone cannot replicate.

However, the closer one gets to implementation, the blurrier the picture becomes as potential use cases proliferate. Given the breadth of central bank mandates and the range of plausible use cases, the challenge will be finding a disciplined framework to assess the value appropriately. Essentially, the broader the central bank mandate is, the easier it is to see a wide range of AI applications on the horizon, but the harder it becomes to identify, up close, which will deliver the most meaningful value. This is the farsightedness paradox the central banks face.

Exhibit 1: Central bank lifecycle



Source: Oliver Wyman analysis

© Oliver Wyman

	Lifecycle	Where AI can create value	Typical applications
1	**Monetary policy and macroeconomic analysis**	Improve timeliness, signal extraction, and scenario analysis in support of policy judgement	Nowcasting, alternative-data analysis, scenario generation, research support
2	**Financial stability and macroprudential policy**	Identify vulnerabilities, contagion channels, and correlated behaviors earlier and across more data sources	Network analytics, market-intelligence synthesis, cyber-contagion stress exercises, concentration mapping
3	**Prudential, conduct and market supervision**	Expand coverage, support scarce supervisory capacity, and detect emerging risks more consistently	Complaints analysis, licensing support, anomaly detection, liquidity forecasting, market intelligence
4	**Payments and market infrastructure oversight**	Detect anomalies, cyber signals, and operational weaknesses faster in systems where resilience is time-critical	Real-time payment anomaly detection, fraud/AML alert triage, resilience monitoring
T	**Shared services and institutional operations**	Improve the internal operating engine that enables the central bank mandate	HR service support, budgeting analytics, procurement support, IT helpdesk chatbot

T = Transversal

Across the central bank lifecycle, four activities stand out as particularly promising in terms of concrete institutional value:

Macroeconomic forecasting

AI ensembles and alternative data sources (such as billing flows, card transactions, mobility indicators, and web activity) can improve nowcasting and enable earlier detection of turning points, helping shorten policymaker reaction times.

Risk-based supervision

AI techniques, such as anomaly detection and natural language processing, can expand supervisory coverage across large populations of firms. Applications include automated review of supervisory returns, detection of suspicious trading patterns across counterparties, and early warning signals for contagion risks in interbank networks.

Licensing application review

AI can assist with the initial assessment of licensing applications by extracting relevant information, checking submissions against predefined criteria, and highlighting areas requiring supervisory judgment, thereby accelerating reviews and improving consistency across cases.

Payments and infrastructure surveillance

Streaming machine learning models applied to transaction data can support near-real-time detection of fraud, AML risks, and cyber threats, thereby strengthening payment system resilience.

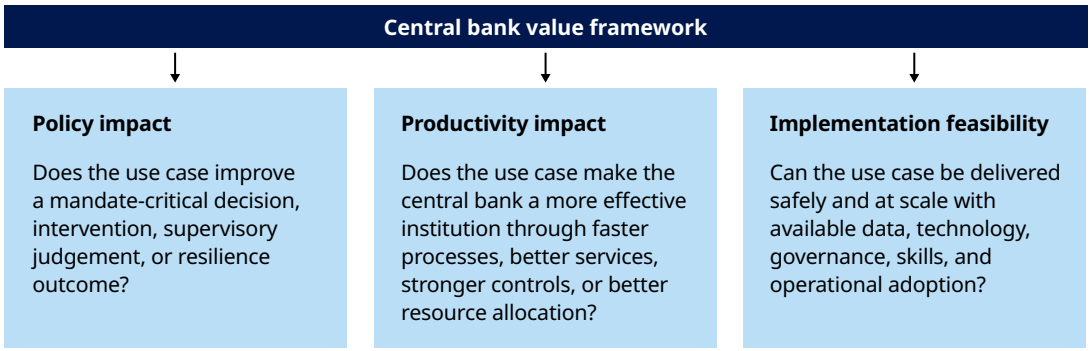
Consequently, central banks need a way to distinguish applications that are simply interesting from those that improve policy judgment or institutional effectiveness in a sufficiently measurable way to justify investment.

© Oliver Wyman

A practical framework for central bank AI prioritization

The private sector often prioritizes its AI use cases based on revenue uplift or cost reduction. Central banks cannot. The relevant question is whether AI changes the quality, speed, coverage, resilience, or explainability of an activity within the value chain.

A practical central-bank value framework should evaluate AI use cases by their impact.



Assessing **the policy impact** of AI use cases requires careful judgment. Unlike productivity gains, which can often be quantified through operational metrics, policy impact is frequently indirect, preventive, and closely linked to public outcomes, making it inherently more difficult to measure.

In this regard, two considerations are particularly important. First, measurement culture may need to evolve. Central banks are often less accustomed than private-sector organizations to working with imperfect or proxy-based value metrics. Yet many of the most important AI benefits cannot be measured with precision and will require a greater institutional willingness to act on directional evidence.

Second, policy impact is often harder to assess but more critical for the central bank mandate. In practice, the most strategically important applications are often the hardest to implement, while the easiest to deploy may offer more limited policy value. One way to approximate policy impact is whether the application can be linked to a concrete improvement with respect to the central bank mandate (see exhibit 2 for sample questions).

© Oliver Wyman

Exhibit 2: Sample questions to assess policy impact

	Lifecycle	Key questions to assess policy impact of a possible application
1	**Monetary policy and macroeconomic analysis**	Would this help the policy committee detect an inflection point earlier or understand uncertainty better?
2	**Financial stability and macroprudential policy**	Would this reveal a vulnerability, a concentration, or a contagion path that current monitoring may miss or only pick up with material delay?
3	**Prudential, conduct and market supervision**	Would this allow supervisors to intervene earlier, cover more firms while staying proportional or challenge riskier behaviour more consistently?
4	**Payments and market infrastructure oversight**	Would this reduce the likelihood or impact of disruption in a time-critical infrastructure?

Source: Oliver Wyman analysis

On the productivity impact side, value is often more straightforward to quantify. It can stem from:

- Staff hours saved in research, analysis, or reporting
- Faster turnaround times for automatable processes
- Reduced error rates and improved consistency
- Greater coverage achieved with the same or only marginally greater resources
- Higher adoption of tools that improve workflow efficiency over time

Implementation feasibility is the third lens of the framework and assesses whether a high-value AI use case can realistically be delivered at scale. While AI capabilities continue to advance rapidly, implementation is often constrained less by the technology itself than by organizational readiness, infrastructure, and delivery capacity. In this regard, four considerations are particularly important:

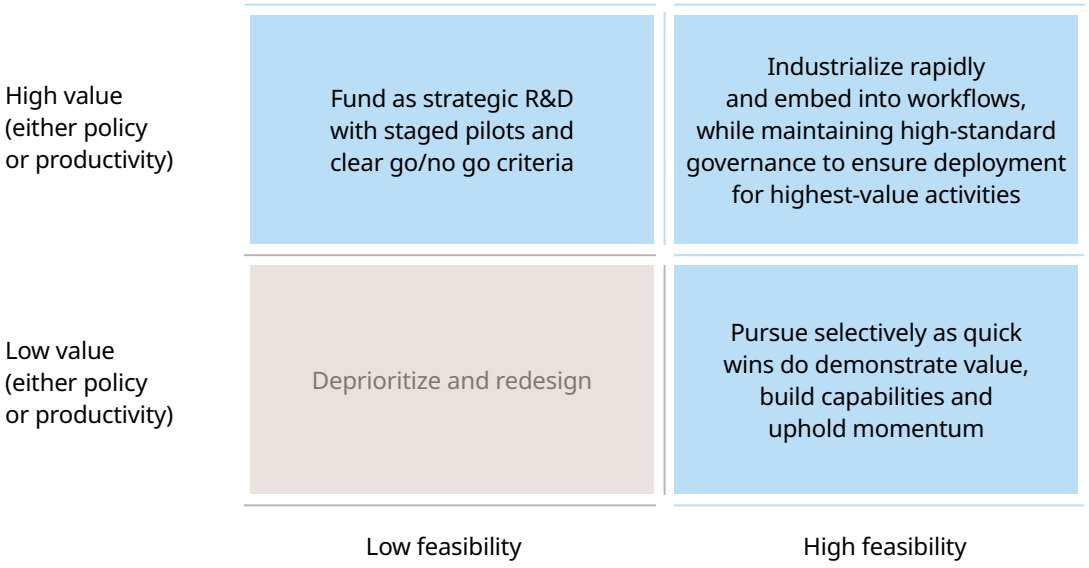
- **AI operating costs can escalate quickly**. As AI adoption scales, token consumption and inference costs can rise rapidly, particularly with more sophisticated reasoning models. Central banks should ensure that increasing AI expenditure is matched by commensurate policy or productivity gains.
- **Institutional culture may slow adoption**. Central banks are built on deep subject-matter expertise, rigorous analysis, and a naturally risk-averse culture. Preparing analysis or briefing senior decision-makers is often seen as a core professional responsibility rather than a task to be delegated to AI. Successful adoption therefore requires careful change management and sustained confidence in AI-supported ways of working.
- **Central banks' technology constraints are materially different from those in the private sector**. Requirements around sovereignty, independence, security, and resilience can significantly constrain deployment options. At the same time, keeping pace with frontier AI capabilities can require substantial investment in computing infrastructure, data platforms, and scarce specialist talent. Remaining at the cutting edge may simply not be economically viable for many institutions.

- **Delivery capacity is inherently constrained**. AI initiatives compete with broader technology, cyber, payments, and SupTech transformation programs for scarce funding and skilled resources. Prioritization should therefore reflect both technical feasibility and the institution's capacity to deliver change.

While the three lenses provide a structured way to assess individual AI applications, they do not automatically determine which initiatives should be pursued first. Prioritization across the portfolio remains challenging for several reasons.

- **First, central banks need an AI use case portfolio that balances policy impact and productivity impact**. Central banks do not operate on the basis of revenue or profitability, but on stability, credibility, and policy effectiveness. As a result, productivity-oriented applications often secure faster senior buy-in because their benefits are immediate and measurable, whereas policy-oriented applications may be strategically more important but harder to quantify.
- **Next, societal value and institutional value should be distinguished**. AI may generate benefits both for the institution itself and for the broader public ecosystem, including citizens, corporates, and market participants. A robust value framework should capture both dimensions.
- **Finally, impact and feasibility should not be treated as a static trade-off**. Over time, they should reinforce one another. Early deployment of more feasible applications can build data assets, delivery capabilities, governance maturity, and institutional confidence, making more ambitious, higher-impact applications easier to pursue. In that sense, prioritization should function less as a one-off filter and more as a flywheel: Practical wins strengthen the enabling foundations for strategic applications, and strategic applications in turn create the case for further investment in foundational capabilities.

Exhibit 3: Prioritization matrix for AI use cases



Source: Oliver Wyman analysis

© Oliver Wyman

Five imperatives for embedding AI

AI implementation at scale requires more than isolated use cases. It requires AI to be embedded into the central bank's strategy and operating model. Doing so depends on a set of imperatives that enable AI to be deployed safely and consistently across the institution. To move from initially short-sighted pilots to sustained, farsighted value, central banks need to strengthen several foundational building blocks. We see five imperatives that central banks should consider.

IMPERATIVE 1

Establish central governance and clear accountability

Challenge: Without coordination, AI initiatives risk fragmentation, duplication, and inconsistent standards across the central bank and the wider market.

Response: Internally, central banks should create a central governance model (whether a dedicated AI office, steering committee, or federated hub-and-spoke structure) to set standards, prioritize investments, allocate funding, and drive reuse across central bank functions. Clear ownership, policies, and escalation mechanisms should ensure alignment with institutional objectives and risk appetite, and create a clear internal function that employees can turn to when they encounter challenges with individual applications and use cases.

Externally, central banks should help steer the responsible adoption of AI across the financial sector through appropriate supervisory expectations and regulatory guidance, helping to support financial sector stability and resilience. We emphasize that this mandate cannot be addressed by micromanaging what every financial sector participant or supervised institution does. It has to be tackled from the privileged farsighted perspective and should be principles and systems-based, rather than use-case-focused or technology-focused.

IMPERATIVE 2

Embed strong risk, ethics, and responsible AI controls

Challenge: Governance principles alone do not ensure responsible AI use. Without practical controls and institutional capability, AI risks remain difficult to detect and manage.

Response: Model risk management should be extended to cover AI-specific risks, including hallucinations, model drift, feedback loops, and adversarial threats. Independent validation, continuous monitoring, auditability, explainability, and human oversight must be embedded throughout the AI lifecycle.

Specifically, central banks need to move beyond generic AI awareness toward role-specific training, practical exercises based on current institutional challenges, and dedicated AI

© Oliver Wyman

assurance capabilities that prepare staff to reliably recognize hallucinations, identify non-auditable or inappropriate AI outputs, and apply informed judgment rather than defaulting to AI recommendations (“cognitive surrender”).

IMPERATIVE 3

Build shared data, technology, and experimentation capabilities

Challenge: Poor data quality, fragmented technology, and legacy infrastructure make it difficult to translate AI ambition into reliable institutional value.

Response: Central banks should use the momentum created by early AI adoption to modernize the data and technology foundations that enable future scale. They should also invest in common data standards, shared data environments, ModelOps capabilities, and reusable AI platforms across the institution.

AI itself can accelerate data cleansing, standardization, and metadata enrichment at a scale that was previously impractical, making this a unique opportunity to improve the quality of the underlying information estate while enabling future AI applications.

IMPERATIVE 4

Invest in talent, delivery capacity, and external partnerships

Challenge: Scaling AI is as much a people challenge as a technology challenge.

Response: A key to tackling the people challenge is building multidisciplinary teams that combine policy, supervision, engineering, product management, and change delivery expertise. The goal is to strengthen internal capabilities through targeted hiring, training, and rotational programs, while leveraging academia, secondees, and carefully governed external partners to access scarce specialist expertise.

IMPERATIVE 5

Start with a small number of priority applications

Challenge: The range of potential AI applications across broad central bank mandates can be overwhelming, making it difficult to determine where to start.

Response: Central banks must avoid trying to address the entire AI opportunity at once. Instead, they need to translate long-term ambition into a small number of high-impact applications that deliver tangible organizational and policy value today, while building the capabilities and confidence needed for broader adoption over time.

Central banks also need to take on the AI challenge early to steer clear of the institutional paralysis that can afflict organizations when they get overwhelmed by the number and variety of AI use case options and applications. Making early progress in establishing priorities should help central banks build lasting momentum institution-wide. The experience gained should in turn refine the framework that helps set priorities, strengthening the metrics, assumptions, and value drivers used to identify the next wave of AI opportunities.

© Oliver Wyman

Ultimately, avoiding paralysis requires balancing long-term ambition with deliberate near-term action. A clear institutional AI ambition should serve as the central bank's North Star, while the imperatives above translate that ambition into practical capabilities and sustained organizational progress. Establishing this ambition also requires making explicit choices about the role AI should play in institutional decision-making. Central banks, more than most organizations, will be forced to define clearly where AI should inform judgment, where it may augment decision-making, and where accountability must remain unequivocally human.

Quantum computing: Tackling a future central bank imperative today

For central banks, the significance of quantum computing lies in the transformative risks and opportunities it may introduce across financial stability, cybersecurity, payments infrastructure, monetary operations, and regulatory oversight over the next decade.

Beyond the inherent risks and opportunities brought by quantum computing, the intersection of quantum computing and AI is particularly relevant:

- Enhanced financial stability and systemic risk analysis, with AI identifying patterns and vulnerabilities and quantum computing enabling more realistic modeling of highly interconnected financial systems.
- Next-generation economic forecasting through quantum acceleration of machine-learning workloads and AI's ability to design a broad range of policy scenarios.
- Payments and financial market infrastructure monitoring, with AI identifying anomalies and quantum computing improving optimization.
- Quantum-enhanced foreign reserve portfolio optimization with dynamic rebalancing and analytics provided by AI tools.

© Oliver Wyman

Conclusion

AI offers central banks a strategic opportunity to strengthen supervision, improve policy effectiveness, and enhance institutional resilience. Realizing that opportunity, however, requires more than isolated experimentation. It requires AI to be integrated into institutional strategy and embedded within the operating model.

Success will depend on central banks' ability to prioritize use cases in which AI creates the greatest value, builds needed capabilities for safe deployment at scale, and recognizes that some important benefits will be strategic and preventive rather than immediately measurable. In an increasingly AI-driven and cyber-exposed financial system, central banks that act early to embed AI into the way they operate will be better positioned to preserve credibility, effectiveness, and oversight relevance.

Lastly, quantum computing will continue to drive technological disruption even as AI has yet to be fully absorbed. Leveraging lessons learned and becoming a responsive, agile organization will remain essential to keeping pace with technological change, ensuring that central banks operate with long-term impact and vision in mind.

© Oliver Wyman

Oliver Wyman, a Marsh (NYSE: MRSH) business, is a management consulting firm driven by deep industry insight, bold innovation, and a collaborative approach that cuts through complexity to help organizations navigate their most defining transformative moments. Marsh is a global leader in risk, reinsurance and capital, people and investments, and management consulting, advising clients in 130 countries. With annual revenue of $27 billion and more than 95,000 colleagues, Marsh helps build the confidence to thrive through the power of perspective.

For more information, visit oliverwyman.com, or follow us on LinkedIn and X.

Americas	Europe	Asia Pacific	India, Middle East & Africa
+1 212 541 8100	+44 20 7333 8333	+65 6510 9700	+971 (0) 4 425 7000

AUTHORS

Dominik Weh
Partner, Global Head of Central Banks and Financial Authorities Platform
dominik.weh@oliverwyman.com

Alexander Ovsepyan
Global Chief of Staff of Central Banks and Financial Authorities Platform
alexander.ovsepyan@oliverwyman.com

Louis Bourbouze
Engagement Manager, IMEA Lead of Financial Innovation, Central Banks and Financial Authorities Platform
louis.bourbouze@oliverwyman.com

CONTRIBUTORS

Davide Taliente, Partner, Global Chair of Government and Public Institutions

Jad Haddad, Partner, Global Head of Quotient, AI by Oliver Wyman

Yaroslav Zakharov, Principal, Central Banks and Financial Authorities Platform, Middle East and Central Asia

Ugur Koyluoglu, Partner, Head of Government and Public Institutions, Americas

James Twaddle, Partner, Head of Government and Public Institutions, Australia

Maria Fernandes, Partner, Central Banks and Financial Authorities Platform, Europe

Martin Stolz, Partner, Central Banks and Financial Authorities Platform, Europe

Leticia Rubira, Principal, Central Banks and Financial Authorities Platform, Europe

Copyright 2026 Oliver Wyman

All rights reserved. This report may not be reproduced or redistributed, in whole or in part, without the written permission of Oliver Wyman and Oliver Wyman accepts no liability whatsoever for the actions of third parties in this respect.

The information and opinions in this report were prepared by Oliver Wyman. This report is not investment advice and should not be relied on for such advice or as a substitute for consultation with professional accountants, tax, legal or financial advisors. Oliver Wyman has made every effort to use reliable, up-to-date and comprehensive information and analysis, but all information is provided without warranty of any kind, express or implied. Oliver Wyman disclaims any responsibility to update the information or conclusions in this report. Oliver Wyman accepts no liability for any loss arising from any action taken or refrained from as a result of information contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages. The report is not an offer to buy or sell securities or a solicitation of an offer to buy or sell securities. This report may not be sold without the written consent of Oliver Wyman.