

SCALING AI IN THE EU FINANCIAL SECTOR



OLIVIER FLICHE

Director, Innovation, Data
and Technological Risks
Directorate – Autorité de
Contrôle Prudentiel et
de Résolution (ACPR)

From adoption to supervision: building trustworthy AI in the financial sector

Four years apart, in 2021 and 2025, the ACPR surveyed the main players in France's banking and insurance sectors regarding their adoption of artificial intelligence (AI). The finding was clear: whereas in 2021 AI remained largely an experimental field for many firms, it has now been adopted by almost all of the institutions surveyed.

This adoption concerns both "traditional" AI (machine learning) and generative AI, which emerged more recently but has spread more rapidly. The primary drivers are the expectation of operational efficiency gains through the automation of time-consuming processes and opportunities to improve the customer experience, which is an even more important factor in the adoption of generative AI technologies.

To a lesser extent, institutions also mention opportunities to strengthen risk management and generate new sources of revenue.

With regard to AI development strategies, all institutions have adopted a hybrid approach, combining in-house development, procurement of off-the-shelf solutions, and co-development with service providers depending on the specific use cases involved. Large Language Models (LLMs) are overwhelmingly purchased as off-the-shelf solutions. Financial-sector firms are now also beginning to explore the potential of agentic AI.

The key risks associated with this widespread adoption should not be underestimated: systemic risks arising from dependence on a small number of technology providers, heightened cyber threats, risks of loss of control by institutions, as well as risks of discrimination, privacy infringements, and customer manipulation. The opacity of algorithms, often referred to as the "black box" problem, represents a major concern both for customer protection and for financial institutions' ability to manage risks effectively. At present, the regulatory framework rests on two pillars: the European AI Act, which notably addresses two high-risk financial use cases, and the existing banking and insurance regulatory framework, which is technology-neutral and applies to AI in the same way as to any other process or tool used to provide financial services. The ACPR will become a Market Surveillance Authority under the AI Act as of 2 December 2027.

**Maintaining a close
dialogue with industry
while also strengthening
cooperation among
Authorities.**

The ACPR's supervisory approach to AI is risk-based and aims to identify the most risky AI systems and, where appropriate, subject them to in-depth technical review. As an initial step, the ACPR is developing assessment methodologies covering governance, robustness, fairness, explainability, and cybersecurity.

To support this work, the ACPR has developed a co-construction approach that allows it to draw both on industry best practices, through technical workshops involving volunteer institutions, and on ongoing academic research. In this context, on 1 July it published, for public consultation, a discussion paper on algorithmic fairness in the financial sector. The objective of this paper is to shed light on one of the main requirements applicable to high-risk AI systems under the AI Act, namely the obligation not to discriminate, in a sector where risk-based pricing segmentation is deeply embedded in risk management practices. The paper seeks to clarify the key issues and debates by bringing together legal requirements, data science techniques, and practices observed across the sector, and to provide financial stakeholders with practical benchmarks to help them move from reflection to implementation.

The AI Act is the cornerstone of AI regulation in Europe, but it may not represent the final or sole regulatory destination. The growing risks that AI poses to cybersecurity and sovereignty have prompted discussions on both sides of the Atlantic about new, complementary forms of AI regulation and oversight. In the financial sector alone, new questions are already emerging. For example, how can a payment transaction initiated and executed by an AI agent be secured in the future? Similarly, should investment recommendations generated by general-purpose AI assistants eventually be regulated, even if such tools do not currently fall within the scope of financial or insurance intermediation rules?

Building "trustworthy AI" in the financial sector is clearly no easy task, given the rapidly evolving nature of the technology. However, the regulatory foundations are now in place in Europe. It is therefore incumbent upon supervisory authorities to take ownership of this issue. Given the complexity of the challenges involved, authorities stand to benefit greatly from maintaining a close dialogue with industry while also strengthening cooperation among themselves at sectoral and cross-sectoral levels, both within Europe and internationally.



MARIUS SKUODIS

Board Member –
Bank of Lithuania

Why AI alone will not transform European finance?

Imagine the following scenario: One morning, you open your banking app to discover that your AI financial assistant has spent the night comparing every mortgage offer on the market, negotiated better terms with another lender and prepared all the necessary paperwork. It now recommends that you refinance, explaining why the new offer is superior and asking whether you would like to proceed.

Science fiction? Maybe. This scenario, however, shows how AI could become genuinely transformative. Today, many Europeans hold mortgages, insurance policies or investment products that are not necessarily the best fit for their needs. While better alternatives may exist, identifying, comparing and evaluating such choices takes time, effort and expertise. AI agents could reduce these costs to almost zero.

In a nutshell, this illustrates AI's potential to fundamentally reshape competition in financial services. Consumers would no longer need to monitor the market or possess deep financial expertise. Their AI assistant could continuously search for better products, negotiate improved terms and identify opportunities that would otherwise remain unnoticed.

Nonetheless, it would be a mistake to assume that technology alone will transform finance. **Technological breakthroughs deliver their greatest benefits when supported by the right institutions and accompanied by appropriate management of emerging risks. AI is no exception.**

Consider the institutional environment. Imagine that your AI identifies the best mortgage for you – not in your home country but elsewhere in the European Union. Can you seamlessly switch? The answer is no. While home-country switching can be costly and complicated, receiving financial services across borders entails still higher levels of expense and complexity. Differences in insolvency regimes, contract law, taxation and consumer protection continue to fragment Europe's financial markets.

The more effectively AI matches consumers with the best financial products across Europe, the greater the economic cost of the remaining barriers that prevent those opportunities from being realised. **By reducing information frictions, AI raises urgent need to return to the EU roots – European economic integration.** Completing the Single Market – particularly in financial services – thus becomes not merely a longstanding European objective but a necessary complement to Europe's AI ambitions. AI can identify enhanced opportunities but only an integrated market can ensure that consumers have access to them.

AI raises urgent need to return to the EU roots – European economic integration.

Managing AI-related risks will be equally important. The same AI capabilities that allow financial institutions to detect fraud or manage anti-money laundering and counter-terrorist financing risks can also be exploited by criminals to hide the movement of illicit funds or create convincing scams, impersonation attacks and cyber-enabled financial crime.

Yet perhaps the most fundamental risk lies elsewhere. A resilient financial system cannot rely exclusively on technology – regardless of how resourceful that technology becomes. Financial supervision ultimately requires judgement, accountability and the ability to respond to events that fall outside historical experience. AI is a good tool for information processing

and otherwise supporting supervisory judgement. However, it cannot bear responsibility for preserving financial stability. Especially during periods of financial stress, supervisory judgement cannot be outsourced to algorithms.

Last but not least, there is the question of how much these new AI tools will cost. AI will not become transformative simply because better models become available. Unlocking its potential requires substantial investment in data, computing infrastructure, cybersecurity, organisational change and, above all, people capable of deploying these tools effectively.

For these reasons, investing in AI must go hand in hand with investing in people to ensure the ultimate layer of resilience. Supervisors need the expertise to understand AI, challenge its outputs and intervene when models fail or circumstances change.

In summary, AI may be one of the most powerful technologies to reshape European finance in decades. Ultimately, however, its success hinges on more than the sophistication of its algorithms. Completing the Single Market, investing in data and people, and ensuring that AI remains a tool that strengthens – rather than replaces – human judgement will determine whether Europe merely adopts AI or truly harnesses its transformative potential.



DAMIAN JAWORSKI

Executive Director – European Insurance and Occupational Pensions Authority (EIOPA)

Scaling AI in finance: systemic risks, resilience and European sovereignty

Artificial Intelligence (AI) is moving from experimentation towards established use across financial services in general, including in the insurance sector. EIOPA's 2025 Generative AI market survey, covering 347 insurers in 25 EEA countries, found that 65% were already using generative AI and another 23% planned to do so.

As AI is deployed across the insurance value chain, it creates both opportunities and new financial stability challenges. AI technologies are increasingly used in fraud detection, customer service, claims management and data analysis.

These technologies may improve efficiency and support innovation; however, they also amplify cyber and operational risks, increase reliance on external technology providers and may contribute to market concentration and interconnectedness.

Advanced AI systems may also increase the sophistication of cyberattacks, automated fraud and social engineering techniques, making cyber risk more

dynamic and potentially more difficult to model using traditional actuarial approaches.

Against this background, EIOPA, alongside the other two European Supervisory Authorities, has raised awareness about the ICT risks posed by the widespread adoption of advanced AI models and engaged with EU competent authorities to ensure that financial entities take appropriate mitigation measures. In their first annual report on major ICT-related incidents under DORA, the ESAs encouraged financial entities to strengthen cybersecurity measures to maintain their resilience amid the rapid evolution of highly capable AI-driven tools.

The ESAs also concur with the ESRB's recent warning on the systemic cyber risks stemming from frontier AI models, whose ability to identify and exploit high-severity vulnerabilities within very short timeframes has advanced significantly. Financial entities are encouraged to make appropriate arrangements to adapt their cybersecurity capabilities, and competent authorities to reflect these developments in their supervisory activities. They can also strengthen defence through faster detection and response.

These developments underline the importance of strong governance frameworks, effective risk management practices and robust supervisory oversight. They also expose Europe's dependence on non-European providers for critical IT infrastructure and AI technologies.

The safe scaling of AI requires effective supervision alongside a diversified European AI ecosystem.

A strong framework, with implementation as the priority

The EU already has a strong regulatory framework. DORA addresses ICT risk management, testing, incident reporting, contractual safeguards, concentration risk and exit strategies, and establishes oversight of critical ICT third-party providers. The AI Act complements it through requirements for AI systems and GPAI model providers, including additional obligations for models presenting systemic risks. In EIOPA's survey, 59% of respondents considered the AI Act's generative AI provisions

useful or very useful for improving provider reliability.

These instruments do not specifically identify correlated model behaviour, common dependencies across models and integrators, or simultaneous AI-driven market reactions, but they enable it. Therefore, effective implementation and coordinated supervision are more effective than another legislative layer. Financial entities and supervisors should map material dependencies, test provider or model failures, verify that exit plans work, and improve information-sharing. Financial entities should retain human expertise, fallback capacity and sufficient provider diversity.

CADA and the supply side

The Cloud and AI Development Act (CADA), the AI Continent agenda and the European Technological Sovereignty Package can complement these efforts by addressing the supply side. CADA aims to triple the EU's data-centre capacity within five to seven years, accelerate infrastructure deployment and establish an EU-wide framework for assessing cloud and AI sovereignty. Together with AI Factories, AI Gigafactories, investment support, skills initiatives and the EU Open-Source Strategy, these initiatives could broaden access to European computing infrastructure and AI models while fostering a more diverse and competitive ecosystem.

However, while sovereignty enhances resilience, the two are not synonymous. Reducing dependence on a handful of non-EU providers only to create dependence on a handful of EU providers would leave the underlying concentration risk largely unchanged. Sustained efforts to promote competition, interoperability and portability therefore remain essential to building long-term resilience.

To conclude, AI scaling is beneficial and should be enabled, but its dependencies must be visible and manageable. At this stage, the EU does not need another broad layer of AI regulation for finance. The priority is rigorous implementation of DORA and the AI Act, coordinated monitoring and targeted supervisory action wherever concrete gaps emerge.



STEPHEN MARTIN

Head of FSI Security and Compliance EMEA & APJ – Amazon Web Services (AWS)

Scaling AI in EU finance: the technology and policy conditions for success

Artificial intelligence is already transforming Europe's financial sector enhancing fraud detection, improving financial customers experience with personalisation, modernizing payments, and addressing legacy systems supporting legacy refactoring. The financial sector is using AI to streamline compliance, drive efficiencies and create a focus on customer experience and AWS is supporting this by providing access to the latest and most innovative AI models and services to its customers. AWS supports customers with the highest levels of security, data protection and resilience for their business services.

To do this there are four areas that our customers tell us are key.

The first is cloud infrastructure. Cloud-based AI enables institutions of every size — from fintech start-ups to global banks — reach sophisticated capabilities without building costly proprietary systems. A layered approach using purpose-built infrastructure to train models and run inference; tools and foundation models to build applications;

and ready-to-use applications requiring no machine-learning expertise levels the competitive field and accelerates innovation.

The second is data management. Data is what turns generic models into valuable applications. Institutions need access high-quality datasets, robust data-quality controls, and the ability to comply with data-residency rules, while allowing for the data to be moved across geographies where appropriate. Moreover, access to varied data reduces the risk of correlated models and herding behaviour, allowing for institutions drawing on different data to reach different conclusions. Data-localisation requirements and regulatory complexity can restrict access to the data, computing and cross-border collaboration that supports the innovation that AI depends on.

The third is model deployment, which rests on the shared responsibility model. AWS embeds significant security, safety and responsible-AI controls at the infrastructure and model layers; financial institutions add controls specific to their use case and regulatory context. This includes model evaluation, configurable guardrails, bias detection, and transparency documentation that set out intended uses and limitations. Institutions remain accountable as sophisticated and informed users who understand, test and monitor the technology they deploy.

The fourth is cybersecurity and resilience. AI is both a threat vector and a defensive capability: it can increase the ease and impact of cyberattacks, while materially strengthening detection, response and recovery. The practical implication is important — declining to adopt AI in cyber defence means facing AI-enabled adversaries without AI-enabled defences. For emerging agentic AI, security is most reliable when enforced through deterministic controls external to the agent — identity and access management, scoped permissions, complete audit trails and external override — rather than the model's own reasoning.

How can technology providers support responsible adoption at scale? By supplying the tools and guidance that translate responsible AI from theory to practice, by helping close the digital-skills gap that 61% of businesses say slows their performance, and by supporting sandboxes and public-private collaboration where risks are not yet fully understood.

Is the EU framework sufficiently clear, proportionate and innovation-friendly? The AI Act provides valuable risk-based guardrails, and its ambition to protect

rights while enabling innovation is welcome. However, AI is already covered by a broad body of law. The focus should be on high-risk applications rather than duplicating existing requirements or layering horizontal regulation onto sectoral rules. Where existing law falls short, we would recommend amending it rather than building a parallel regime.

Against other jurisdictions, the decisive factor is international coordination. Fragmented, conflicting requirements restrict access to data, computing and talent, and disproportionately burden smaller firms. Consistent, interoperable standards — technology-neutral, outcomes-focused and proportionate — will let Europe's financial sector innovate responsibly and compete globally.

The technology to scale AI in Europe's financial sector exists today. The task now is a policy framework proportionate enough to let it flourish.



EMILY PRINCE

Group Head of
Enterprise AI – LSEG

Scaling AI across financial services: trust, choice and resilience

AI is moving from experimentation to becoming part of the operating fabric of financial services. The opportunity is not just automating tasks but improving how professionals work end to end – from discovery of information, monitoring of risks to decision-making. In a sector where confidence rests on accuracy, resilience and accountability, scale will depend less on the novelty of a model than on the quality of the system around it: trusted data, clear governance, human judgement and infrastructure designed for continuity.

For LSEG, this starts with the combination of deep financial-services expertise, trusted data and critical market infrastructure. Our 'LSEG Everywhere' strategy has three priorities: making licensed, AI-ready data available where customers work; developing differentiated AI-enabled products; and using AI responsibly across our enterprise. Drawing on our financial markets experience, we are putting this into practice by designing AI in line with the requirements of regulated workflows—accurate, traceable data, effective controls, interoperability and choice.

For example, open standards such as Model Context Protocol (MCP) can help

make AI more useful and governable in regulated workflows by allowing models to connect securely and quickly to governed data. This supports more accurate, context-aware outputs while preserving controlled access, auditability and the ability to switch between models, applications or infrastructure providers. Such recognised standards point to a broader transformation: from AI as a stand-alone tool to AI effectively embedded in a live, flexible ecosystem of models, specialised datasets and technology environments that can be adapted to the intent of financial service professionals.

The next phase should be judged by whether AI delivers measurable improvements in real financial-services processes, across risk, compliance and operational outcomes. AI can synthesise information, identify anomalies, support surveillance and strengthen risk monitoring. However, decisions should not be delegated merely because a model is powerful. Model outputs may be probabilistic or difficult to validate, while accountability remains with regulated firms and their people.

**Diversification and
choice are key to avoid
overdependence, whether
it is defined at the level
of an individual company
or a jurisdiction.**

Adoption therefore requires the right system for the right task: deterministic controls where repeatability is essential and probabilistic models where uncertainty is intrinsic, explainability and assurance proportionate to risk, meaningful human oversight, and tested fallbacks. Model diversity and modular architecture can reduce concentration risk, common-mode failures, and single-provider dependence, if they do not obscure accountability or create unmanaged complexity.

Diversification and choice are key to avoid overdependence, whether it is defined at the level of an individual company or a jurisdiction. Europe is right to invest in data centres, domestic cloud and computing power, skills and AI capability. It is key to support choice and ensure operational resilience and we believe the proposed Cloud and AI Development Act and the wider AI Continent agenda can expand European capacity. Sovereignty should mean strengthening resilience and competitiveness. But requirements that narrow access to the best

available infrastructure could increase concentration in a small number of firms, raise costs, and make European firms less resilient and competitive.

Instead, firms should be required to operate based on principles of 'geopolitical resilience' and effective control: firms should understand critical dependencies, move data and workloads securely, test exit and recovery plans, and retain meaningful choice across models and cloud providers.

The EU already has a substantial framework. The AI Act governs AI-specific risks, while DORA addresses operational resilience and critical ICT dependencies; existing financial-services rules continue to apply to governance, outsourcing and accountability. The immediate priority should therefore be coherent implementation. Supervisors and the AI Office should clarify how horizontal and sectoral rules interact, coordinate interpretations across the EU and recognise existing financial-sector controls. International alignment is equally important because fragmented standards or restrictions on cross-border data flows would raise costs and weaken the evidence base on which reliable AI depends.

Europe can combine ambition with trust if it treats openness and resilience as complementary. Policy should remain technology-neutral and outcome-focused. Regulation should encourage diversification, portability, interoperability and multi-provider strategies to ensure 'geopolitical resilience' rather than prescribe provider location.

The goal should not be technological autonomy at any cost, but strategic capability with choice: an environment in which European financial services can scale AI responsibly, adopt the best technology when appropriate, switch when necessary and continue serving markets through disruption.



MICHAL PAPROCKI

Chief Technology Officer
– Euroclear Group

AI in post-trade: from experimentation to trusted transformation

Technology and FMI

Technology has always shaped financial services. From dematerialisation and real-time messaging to cloud, and distributed ledger technology, each new wave has changed how markets connect, settle and manage risk. For a financial market infrastructure, technology is never an end in itself. It must strengthen trust, resilience and efficiency at scale.

At Euroclear, this is how we approach artificial intelligence. We see it as a strategic capability that can help capital markets become more efficient and resilient, provided it is deployed with the same discipline that underpins our core infrastructure. Our ambition is not to add AI at the margins, but to become an AI-augmented FMI: one where human expertise, high-quality data and trusted technology combine to improve outcomes for clients and the wider market.

Euroclear and AI

This journey did not start with the current wave of generative AI. Euroclear has used advanced analytics and

machine learning for several years. More recently, generative AI has opened new possibilities for knowledge-intensive activities, software engineering and client service. To move from experimentation to industrialisation, we created an AI Centre of Excellence to identify use cases, define implementation standards and ensure that initiatives are assessed through a common governance framework.

The result is a growing portfolio of AI opportunities. More than one hundred potential use cases have been identified, with active initiatives spanning settlement efficiency, collateral optimisation, sanctions and AML processes, client service, finance and commercial engagement. This breadth matters: AI will create value not through isolated pilots, but by being embedded in processes where data, judgement and repeatability meet.

Promising applications

The most promising use cases are those that support the core mission of a post-trade infrastructure. In settlement, predictive analytics can foresee matching failures before they occur, identify their underlying causes, and enable proactive remediation. This is particularly relevant as Europe prepares for T+1 settlement where the time available to prevent fails will shrink materially. In collateral management, AI can help analyse constraints, eligibility, liquidity needs and optimisation opportunities, supporting better allocation decisions and more efficient use of assets.

Our ambition is not to add AI at the margins, but to become an AI-augmented FMI.

Other applications are equally important, even if less visible to the market. In compliance for instance, AI can support sanctions screening and AML controls, provided outputs remain explainable, traceable and subject to human oversight. This is a good example of where AI can strengthen risk management without weakening accountability.

Scaling AI in core processes

The lesson from our experience is clear: scaling AI in capital markets is primarily a question of trust. First, firms need AI-ready data: secure, high-quality, well-governed and accessible within clear boundaries. Second, AI must be

integrated into resilient technology architectures, with appropriate controls and monitoring. Third, governance must be designed in from the outset, including risk classification, cybersecurity, auditability and compliance with existing financial regulation. Fourth, the operating model must remain human-centric. AI can augment judgement and automate repetitive tasks, but accountability must stay with people.

More needs to be done to ensure these conditions are in place across the industry. Policymakers and supervisors have an important role in supporting responsible innovation without creating unnecessary overlaps between horizontal AI rules, financial services regulation, data protection, cyber resilience and outsourcing requirements. Greater supervisory convergence, clearer expectations on explainability and proportionate approaches to lower-risk use cases would help firms scale AI safely. Industry collaboration will also be essential, because many post-trade inefficiencies are shared across the market and cannot be solved by one institution alone.

The way forward

The human journey is just as important as the technology journey. AI will require literacy across the organisation, deeper collaboration, and a culture that is comfortable experimenting while remaining disciplined on risk. The unanswered questions are therefore not only technical. How do we measure the true value of AI in critical market processes? How do we preserve accountability as systems become more autonomous? How do we ensure that European capital markets remain innovative and competitive while staying anchored in trust?

For Euroclear, the answer is to proceed with ambition and prudence. AI can help financial market infrastructures improve efficiency, resilience and risk management. But its success will depend on the quality of our foundations: trusted data, robust governance, resilient technology and empowered people. If we get these foundations right, AI will not simply make post-trade faster; it will help make markets stronger.



DERA MCLOUGHLIN

Partner, Head of Consulting,
Ireland – Forvis Mazars

Scaling AI: investment, regulation and the transformation challenge

AI is increasingly recognised as a transformative force within financial services. Across banking, insurance, and capital markets, AI is moving from experimentation to operational deployment. While the EU has made significant progress, the question remains as to whether EU financial institutions can successfully scale AI to deliver sustainable business value. This challenge sits at the intersection of investment, regulation and organisational capability.

Evidence suggests that AI adoption within European financial institutions is now widespread. The EBA reported in 2025 that 92% of EU banks had deployed AI, most commonly in AML, fraud detection, customer service, credit assessment, and operational automation. Similarly, the ECB has reported that more than 85% of significant institutions under its supervision utilise AI and GenAI.

The insurance sector demonstrates similar progress. EIOPA has reported that half of non-life and nearly one-quarter of life insurers have adopted AI to support underwriting, claims,

customer engagement, pricing optimisation, and actuarial modelling. These findings suggest that the EU is broadly comparable with other advanced economies in the practical deployment of AI across financial services.

Investment

However, a different picture emerges when investment is considered. Whilst European financial institutions are increasingly investing in AI-enabled transformation, the region lacks the concentration of venture capital, hyper-scale technology firms, and AI infrastructure that characterise other regions such as the USA, where private AI investment significantly surpasses EU levels. This disparity is important because value in AI increasingly depends not only on adoption but also on the ability to develop foundational models and access advanced computing capacity.

The distinction between adoption and innovation is critical. European financial institutions are becoming sophisticated consumers of AI, yet comparatively few are creators of frontier AI systems. Consequently, many rely on technologies developed by firms headquartered outside the EU. This dependence raises important questions of technological sovereignty, resilience, and long-term competitiveness.

**EU financial services
scale AI adoption,
but regulation, data
and investment
constrain value.**

Transformation

Among the most significant barriers to scaling AI effectively and responsibly is data. AI depends on high-quality, accessible, and well-governed data. Many financial institutions continue to operate with fragmented data architectures developed over decades of mergers, acquisitions, and regulatory change. While financial institutions possess vast quantities of data, the challenge is often not availability but usability as data may be distributed across systems, stored in incompatible formats, or be subject to differing governance arrangements. Without consistent data quality, lineage, and governance, AI models risk generating inaccurate outputs and unreliable decisions.

Closely linked to the data challenge is the persistence of legacy technology

infrastructure. Many organisations continue to rely on core systems designed long before the emergence of cloud computing, or generative AI. Replacing critical legacy systems remains expensive and complex. Consequently, institutions frequently find themselves attempting to deploy advanced AI capabilities on foundations that were not designed to support them.

Perhaps the most underestimated challenge is organisational change. AI adoption is frequently approached as a technology initiative when, in reality, it is a business transformation programme. Successful scaling requires changes to operating models, decision-making processes, workforce structures, and culture. Those that achieve the greatest value from AI are typically those that align technology investments with broader strategic objectives and foster a culture of experimentation, learning, and innovation.

Regulation

The USA and several Asia-Pacific markets, particularly Singapore, Japan, and South Korea, have moved more aggressively than the EU to support AI and autonomous systems in financial services. Their regulatory approaches have generally emphasised innovation and experimentation, while relying on existing legal frameworks to manage risk. By contrast, the EU has pursued a more structured approach with the AI Act, the world's first AI regulatory framework.

The framework provides a foundation for responsible AI deployment, and it is welcome but further clarification is required to support scaling. Financial institutions need greater certainty on the interaction between the AI Act, the Digital Operational Resilience Act (DORA), CADA and existing sectoral regulations. Greater supervisory consistency, clearer guidance on GenAI, and harmonised expectations would reduce implementation uncertainty while maintaining appropriate safeguards, and ultimately go some of the way to balancing the tension between innovation and regulation.



SEBASTIAN GEHRMANN

Head of Responsible AI, Office of the CTO – Bloomberg

Trust that scales with AI

Financial services is handing larger workloads to AI, and the workforce has already made the shift. In a recent survey of 500 London-based finance professionals conducted by Bloomberg, 88% shared that they use AI tools daily or weekly and 90% report productivity gains. Those gains are showing up as better work rather than less of it: 73% say their overall workload has not fallen, while around half are redirecting freed up time toward higher-value activities like strategy, decision-making, and client relationships. AI is changing work to one of judging and evaluating AI outputs, and the direction of travel is toward delegating more tasks to increasingly agentic systems.

Finance is a natural place to put AI to work, given the volume of data it generates, and one where the cost of error is amplified. The same practitioners adopting the technology are the ones asking for control: in the survey, 51% want human oversight prioritized in risk and compliance and 40% in investment and trading, while only 1% think oversight is unnecessary. Responsible deployment rests on the obligations of everyone involved: buy-side firms acting under fiduciary duties, sell-side firms bound by suitability and supervisory standards, and technology vendors whose systems increasingly support these activities.

At Bloomberg, trustworthiness has guided product development since 1981, long before the current wave of generative AI: a product should be accurate, resilient, and verifiable. Applying that standard to AI has exposed a gap in how content risk is defined. Content risk is whether a system's inputs and outputs meet the safety standards of whoever deploys it, and those standards differ by domain. In finance, safe content is shaped by disclosure obligations, advertising restrictions and suitability standards that generic guardrails know nothing about. That is why our Responsible AI team built and published the first AI content safety taxonomy designed for financial services, organized around where risk actually arises: provenance, where systems may touch confidential data or material non-public information; communication, where client-facing content must stay fair, balanced and accurate; and investment activities, where models may create exposure to manipulation, fraud or inadvertent insider trading (Gehrmann et al., 2025)¹.

We have built purpose-built AI for finance since 2009, treat responsible development as a live discipline, and contribute what we learn back to the research community. For instance, we found that retrieval-augmented generation (RAG), a now-standard way of grounding models in source documents, does not automatically make a system safer and can make it less reliable (An et al., 2025)². The lesson was to engineer around it, grounding answers in trusted data and giving users transparent attribution to verifiable sources. The wider lesson is that safety cannot be certified at the model or technique level or taken on faith from generic benchmarks. AI has to be evaluated in the specific context where a system is deployed.

As finance hands more work to AI, trust must scale with it.

That principle should shape regulation too, and here the ground is uneven. In the U.S., there is still no comprehensive federal AI statute, and the state picture is fragmented and unsettled. For any firm touching the EU market, the EU AI Act is binding, extraterritorial, risk-tiered and increasingly operative in 2026 (though not without its challenges and delays). Across APAC, approaches vary again. Products engineered to one authority's definitions can end up non-

compliant or outdated in another, and the ground may shift before they go live.

The answer is not lighter regulation. It is regulation that travels. Most of the risks AI raises in finance are already regulated: market manipulation, fraud, and insider dealing. AI changes the probability and surface of those violations more than it creates new categories of harm. What this calls for is common, principles-based and technology-neutral standards, with oversight proportional to the risk of each application, built on the sector rules that already exist rather than a parallel regime that dates as fast as this technology does.

Regardless of these challenges, financial services firms are rapidly adopting AI. Doing so in a trustworthy manner is critical for the entire industry. Responsible design choices are integral to our product development process at Bloomberg, and we will continue to adhere to them as we navigate the developing regulatory environments across the globe. Whatever shape the rules ultimately take, Bloomberg remains steadfast in its commitment to responsible AI, and to building products our clients can trust.

1. Gehrmann, S., Huang, C., Teng, X., Yurovski, S., Bhorkar, A., Thomas, N., Doucette, J., Rosenberg, D., Dredze, M., & Rabinowitz, D. (2025). *Understanding and mitigating risks of generative AI in financial services*. In Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency. <https://doi.org/10.1145/3715275.3732168>
2. An, B., Zhang, S., & Dredze, M. (2025). *RAG LLMs are Not Safer: A Safety Analysis of Retrieval-Augmented Generation for Large Language Models*. In Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers). <https://doi.org/10.18653/v1/2025.naacl-long.281>



ELIE FARAH

Managing Partner, Head
of Financial Services for
Europe – Oliver Wyman

Beyond assistants: why scaling AI requires reinventing the bank

Over the past two years, financial institutions have made remarkable progress with artificial intelligence. Thousands of use cases have been identified. AI assistants are moving into customer service, software engineering, research and operations, where many banks are already seeing measurable productivity gains. Large institutions have established AI organizations, governance structures and growing portfolios of production deployments.

Despite that progress, many banks remain focused on the mechanics of adoption rather than what happens next.

Scaling AI is not simply a larger deployment exercise. Moving from 10 use cases to 100 or from one assistant to thousands of employees will not, by itself, produce a different institution. The larger shift comes when AI changes how the bank operates, how decisions are made and where competitive advantage comes from.

Most investment has focused on improving existing operating models. Yet the greater opportunity lies in redesigning them.

In conversations with banking executives over the past year, we've noticed a shift. The discussion has moved beyond copilots and toward questions like whether junior analyst programs still make sense and how relationship managers will supervise AI-generated advice. A year ago, many executives focused on where AI could improve productivity. Today they are thinking about how AI changes organizational design, talent development and competitive positioning. The discussion has become less about deploying tools and more about redesigning institutions around a different economic model.

These leaders see that AI offers much more than improved efficiency across existing processes; it can help a bank redesign client journeys, operations and control functions around AI to improve its entire cost structure and operating model. That allows people to spend less time on routine execution and more time exercising judgment, managing client relationships and supervising increasingly capable AI systems.

Technology alone cannot produce that outcome. The main obstacle is trust.

Customers have always expected banks to make decisions that are understandable, reliable and accountable. That expectation will only become more important as AI systems begin recommending products, comparing providers and, in some cases, acting on behalf of customers. Increasingly those decisions must also be legible to AI agents acting on behalf of customers. Explainability and transparency are becoming commercial capabilities alongside regulatory obligations because both people and machines will influence how financial products are selected and recommended.

**The next competitive
advantage in banking
will come not from
deploying more AI, but
from redesigning.**

The customer relationship is changing. Digital banking strengthened direct connections by making the bank's application the primary interface. AI assistants may weaken that advantage. If customers increasingly rely on intelligent agents to search, compare and execute financial decisions, advantage shifts toward institutions

those agents judge to be trustworthy and reliable.

The workforce presents another challenge. Banking has traditionally developed expertise through apprenticeship, with junior employees learning through repetitive analytical and operational work before taking on greater responsibility. AI is beginning to absorb many of those early-career tasks. Banks will need new ways to develop judgment, client skills and institutional knowledge when much of that training ground has changed. That may require rethinking career paths, performance management and even what qualifies someone to lead teams whose work is increasingly performed alongside AI.

Governance must evolve as well. Banks have spent decades refining governance for people, products and financial risk. AI introduces another operating layer that requires equally clear ownership and accountability. Frameworks designed for dozens of models will not be sufficient when banks manage thousands of AI agents across business functions. Ownership, permissions, accountability, monitoring and lifecycle management become operating disciplines rather than compliance exercises.

Investment may ultimately prove to be the greatest constraint. Building enterprise AI capabilities requires far more than foundation models. Data infrastructure, technology platforms, governance, process redesign, talent and organizational change all compete for capital. Institutions with stronger earnings can reinvest more aggressively, accelerating improvements that weaker competitors struggle to match. That creates a practical advantage: well-capitalized banks can test, learn and redesign faster, while others remain trapped in narrower, incremental programs.

The industry has spent the past two years learning where AI can improve productivity. The next phase will depend on whether banks are willing to rethink how the institution itself is designed. Those that make that transition are likely to build advantages that could shape the competitive landscape for years to come.