

# Scaling AI in finance: key drivers and challenges

---

The chair outlined that the panel discussion would address the drivers of AI adoption, the operational and regulatory challenges hindering its wider uptake, and the adjustments needed in regulatory and supervisory approaches to support deployment.

---

## 1. AI uptake and key drivers in the financial sector

---

### 1.1 Current uptake and emerging trends

An official explained that the UK authorities are jointly conducting regular surveys to assess AI adoption and use cases in order to understand how the financial industry is experimenting with and deploying the technology. The latest results show that 75% of firms already use AI, with a further 10% actively planning to do so, indicating that most financial institutions are now implementing AI in some form.

Current applications are concentrated in the optimisation of internal processes, particularly cybersecurity, fraud detection, process automation, and operational efficiency, areas where firms expect the fastest returns. Over time, adoption is expected to expand into more customer- and market-facing activities, building on progress already achieved in analytical insight and pattern detection to fight fraud and money-laundering.

Another official noted that an extensive industry consultation conducted by the Japanese FSA also indicated that Japanese financial institutions are using AI mainly to reduce internal operational burdens. They remain conservative in customer-facing applications, consistent with broader consumer attitudes in Japan.

An industry speaker expressed the hope that within one to two years, AI will be more integrated into core financial processes such as credit decisions and investment advice, rather than being limited to back-office functions. A key priority going forward will be to measure the return on investment of AI projects more precisely. While many projects are currently funded on expectations of future benefits, firms will need to become more disciplined in applying AI across their organisations and in tracking tangible outcomes, since adoption must ultimately deliver clear business value to justify continued investment.

A public representative anticipated that future discussions will increasingly focus on agentic AI, its use cases and deployment, and how its implementation will require greater coordination across different parts of the financial ecosystem, from payments and banking to digital and data services, as AI models begin to integrate and automate a broader range of activities.

### 1.2 Drivers of adoption

An industry speaker observed that operational efficiency and cost reduction have so far been the primary driver of AI implementation in the financial sector. Looking ahead, two additional factors are expected to gain importance: enhancing products and deepening client personalisation. Asset managers are already using AI to process large volumes of market data, for example to assess the sentiment expressed in companies' financial reports, and to personalise investment solutions that give clients the sense of tailored products while maintaining operational standardisation.

The industry speaker added that successful scaling of AI also depends on user adoption. AI differs from most technologies: it is not an application deployed by a central IT department but a tool that every employee must learn and choose to use. Openness to AI is essential, requiring staff first to understand how AI affects their roles and then to develop the capabilities to use it responsibly.

Another industry speaker agreed that training staff to manage AI-related risks is essential, ensuring that employees understand how to use these tools responsibly and are able to challenge their outputs when needed.

An official added that geopolitical competition is also a significant factor: as with other emerging technologies, countries are racing to establish AI hubs and achieve scale, creating a further incentive for rapid deployment.

---

## 2. Operational challenges and risks from AI

---

### 2.1 Data quality and accessibility

An industry speaker considered that the main challenges to AI adoption today relate to data i.e. data integrity and quality, and determining what data is reliable and usable for AI systems. The availability of adequate, high-quality data should be a core objective of any AI framework, rather than codifying prescriptive rules on how to develop and use AI, which risk turning compliance into a box-ticking exercise that overlooks the essence of responsible use.

Another industry speaker agreed that data quality remains the decisive factor for leveraging AI effectively, since AI only creates value when grounded in reliable, high-quality inputs. Large general-purpose models can perform broad searches but not specialised tasks. Therefore, firms that lack the scale to train new models for specific purposes must augment existing ones with domain-specific data. Concerns over data accessibility have also intensified amid recent geopolitical developments, including questions about using data hosted by US-based hyperscalers.

## 2.2 Hallucination risks

An official emphasised that while AI presents vast potential, its risks must be addressed. A key risk is hallucination, where models generate plausible yet inaccurate information. AI outputs are often so convincing that errors become difficult to detect, unlike with human analysts. This can create consumer-protection risks, for example when robo-advisors provide credible but incorrect investment advice.

The official underlined that because current generative and agentic AI systems rely on models that produce outputs based on probability rather than verified information, the risk of hallucination can never be fully eliminated. Maintaining human oversight remains the most effective safeguard against hallucination, but it is costly, requiring multiple layers of expertise, and firms often lack sufficient incentives to apply it systematically. To ensure the effective implementation of this oversight, three complementary lines of defence were proposed: first, to limit the exposure of retail consumers and investors to hallucination through technical safeguards such as built-in guardrails and systems that verify AI-generated information against trusted data sources (known as retrieval-augmented generation, or RAG); second, to embed human oversight within regulatory frameworks and promote international coordination so that these controls are applied consistently across jurisdictions; and third, to establish safety nets, for instance, insurance products or industry-funded reserves, to protect consumers when erroneous AI-generated advice causes harm.

An industry speaker cautioned against overstating the risks of hallucination, noting that error rates in current large language models are estimated at around 2%, and that comparisons often overlook the fact that humans themselves are far from perfectly accurate.

Another industry speaker compared AI errors to accidents involving self-driving cars, observing that society shows far less tolerance for mistakes made by machines than humans, even though actual error rates are lower. When AI models make flawed decisions, people label them as “hallucinations” whereas similar human mistakes would simply be viewed as incompetence. This difference of perception may hinder the scaling of AI. In addition, when AI systems make recommendations or decisions, their reasoning often lies beyond human comprehension, as it draws on volumes of data no human could process. As a result, errors are difficult to detect. A public representative agreed that discussions on AI hallucination must take into account the potential for human error, reinforcing the need for balanced expectations on both sides.

## 2.3 Impact of AI on human expertise and skills

An industry speaker noted that the erosion of human expertise due to reliance on AI poses another risk: as staff increasingly rely on AI, they may lose the ability to critically assess its outputs, complicating the development of skilled senior professionals.

Another industry speaker stressed that, looking ahead, firms will increasingly compete for individuals capable of challenging how AI thinks, as the focus shifts from implementing the technology to enhancing human understanding and oversight with AI use.

An official agreed that the main discussion in the coming years will increasingly focus on how society adapts to the transformation already underway.

The chair underlined that using AI effectively requires not only willingness, but also the skills and capacity to understand and challenge its outputs. A recent Stanford study also shows that advanced AI models already outperform humans in many professional tasks.

## 3. Regulatory and supervisory approaches

### 3.1 Current regulatory and supervisory approaches to AI

An industry speaker expressed concern over Europe's cautious stance on AI regulation. The US AI Action Plan focuses on practical enablers such as data-centre capacity and power supply, whereas the EU AI Act places greater emphasis on controls and governance. Europe should strike a better balance between protection and innovation. The same applies to data management: while data-protection obligations are necessary, excessive caution over client data and public-cloud use can hinder value creation and scalability. These safeguards are legitimate but need to be reconsidered in light of the goal to foster AI adoption.

The industry speaker also questioned whether regulating the technology itself is the right approach for financial services. Many of the objectives of the AI Act, such as ensuring that AI systems are used appropriately, produce reliable outputs, and operate under adequate oversight, are already pursued in financial services through existing regulatory frameworks and market incentives. Firms have a strong commercial and reputational interest in avoiding poor advice or service, whether it originates from humans or AI, since dissatisfied clients represent a direct business loss.

Another industry speaker highlighted that the current uncertainty surrounding AI regulation is also a significant challenge for the industry. The current phase of AI adoption mirrors the transition to cloud services a decade ago, when it was unclear how national authorities would interpret and apply new rules, with both technologies holding significant potential in terms of efficiency and resilience. This regulatory ambiguity is a tangible barrier to adoption confirmed in a recent survey, where 44% of firms identified unclear regulation as a key obstacle to the implementation of AI systems.

A public representative noted that opinions during the European Parliament's negotiations on the AI Act were divided between those who viewed it as a driver of innovation, providing legal certainty, and those who feared it could constrain adoption through excessive requirements. Both perspectives are valid to a certain extent, and the actual impact will only become clear once the framework is fully implemented next year.

The implementation and supervision of the AI Act also remains challenging: some models will fall under EU-level oversight by the AI Office, but many, particularly in financial services, will be supervised nationally, creating risks of fragmentation. Stronger harmonised supervision is

therefore needed to ensure consistency across the Union. As AI systems become increasingly autonomous and cross-sectoral, especially with the rise of agentic AI, supervision will require a broader and more coordinated approach. The regulatory framework also extends beyond the AI Act to include data and financial-sector legislation, such as the Financial Data Access Regulation (FiDA) and the Data Act, which must be coherent and aligned to avoid overlaps and inconsistencies.

The chair agreed it is still too early to judge the AI Act's impact, before it is fully implemented and emphasised the importance of supervisory convergence across the EU. Guidance is being prepared by EIOPA on the interaction between the AI Act and sectoral insurance rules. Developing such guidance can however be challenging, as it can be perceived as additional requirements rather than practical support.

### 3.2 Adapting regulatory frameworks to a fast-changing AI landscape

A public representative underlined the need for a mindset shift in the way EU AI rules are drafted and implemented, so that they remain effective with the rapid pace of change across AI, data, and digital services. To keep legislation up to date, Europe should move away from a system of fully codified rules towards one of "regulating through supervision and practice", underpinned by clear principles. This already occurs to some extent, as European frameworks are often clarified over time by supervisors to make them applicable or to foster convergence. However, the current process of adjusting Level 1 and 2 texts can take years. Issues could be addressed more effectively through a more flexible approach to regulation and enhanced cooperation.

An official agreed that a responsive and flexible policy approach to AI is essential. The unprecedented pace of technological change demands regulation that remains principles-based, outcomes-focused, and adaptable. Since neither regulators nor firms can anticipate every use case or risk, codifying rules too early would limit the ability to respond effectively to emerging opportunities and challenges. This adaptive, principle-driven philosophy lies at the core of the UK's approach to AI regulation, which builds on existing regimes governing notably market conduct, senior-management responsibility, and operational resilience. The aim is to provide flexibility while maintaining clarity. The FCA's first clarification, issued in April 2024, set expectations for the use of AI within these existing frameworks.

An industry speaker also supported a principles-based approach to AI regulation. A genuine understanding of risks arising from specific use cases is needed before imposing detailed compliance or risk-management measures, as these risks will continue to evolve. Lessons can be drawn from the cloud transition a decade ago, when similar debates arose over how to regulate technological change. A principles-based framework is preferable because it compels regulators and firms to think through the objectives of AI adoption, identify the risks involved, and determine how best to manage them, promoting judgement and accountability over box-ticking compliance.

Another official noted that in Japan discussions on AI regulation for financial services remain at an early stage.

The country's cautious and incremental approach reflects a determination to balance innovation with risk management and to develop rules that both enhance customer confidence and support technological progress. Some form of regulation will eventually be needed to address the specific risks associated with AI use, such as hallucination. In June, a public-private forum was launched to exchange views among stakeholders, enabling authorities to better understand AI use cases and risks and to explore possible directions for future policy that strike an appropriate balance between innovation and effective oversight.

### 3.3 Public-private cooperation and the role of sandboxes

A public representative suggested that strengthening public-private cooperation is essential for developing a common understanding of technological developments and designing relevant regulation. The broader use of regulatory sandboxes, as promoted by the AI Act, one of the first EU laws to make their creation mandatory, should be further developed in order to provide a structured space for this collaboration. Such sandboxes and experimental environments enable public and private actors to test whether rules and models are fit for purpose and to anticipate future challenges, as AI systems become more complex.

An official agreed that the difficulty of predicting the direction and impact of technological disruption calls for ongoing dialogue among stakeholders and the creation of testing environments where regulators and firms can experiment with AI applications safely, gather practical evidence, and adjust regulatory approaches as experience is gained. In the UK, the Digital Regulation Cooperation Forum (DRCF) promotes alignment between government and regulators, coordinating their approaches on cross-cutting digital issues such as AI, online safety, and data governance, to ensure coherent oversight and consistency. Engagement with industry is expanding through initiatives such as the AI Lab and AI Spotlight, which provide collaborative platforms and a use-case library to map AI applications, assess their value, and identify potential policy gaps. Two testing tracks - the Supercharged Sandbox for early model exploration and AI Live Testing for more mature systems - enable firms to trial AI models in real settings, helping regulators identify challenges and clarify expectations on explainability and risk management. While the core objectives of supervision - sound risk management and redress where needed - remain unchanged, supervisors are encouraged to adopt a tech-positive mindset, combining flexibility and collaboration with industry to ensure that AI delivers value for consumers while strengthening the UK's global competitiveness.

An industry speaker welcomed the use of regulatory sandboxes and open exchanges between firms and supervisors, stressing that such cooperation is crucial to share lessons and identify risks that may not be captured by formal regulation.

The chair concurred that sandboxes are a valuable tool to foster experimentation and cooperation, as shown by those that have already been developed in the EU.