

AI and the future of retail financial services

July 2026

Sheldon Mills



Contents

3	Foreword by Sheldon Mills
5	Executive summary
12	The AI autonomy spectrum
14	The systemic driver: Advances in AI capability
23	System shift 1: The transformation of firms
33	System shift 2: New consumer journeys
51	System shift 3: A reshaped competition landscape
68	System shift 4: Amplified financial crime and cyber risk
77	Regulatory implications of AI-enabled finance
93	Priority recommendations for AI-enabled finance
118	Annex I - Glossary
121	Annex II - Engagement paper responses
126	Annex III - Consumer survey findings
130	Annex IV - Consumer focus group findings
131	Annex V - Economic growth impacts
132	Annex VI - International approaches to AI
138	Annex VII - Infrastructure for agentic finance

Foreword

AI is already part of retail financial services. Used well, it can help firms serve customers more effectively, detect harm earlier and support better financial decisions. But as the technology becomes more capable and more widely used, it may also change how consumers engage with financial services, how firms compete and how markets operate.



The FCA Board asked me to conduct this Review into how advances in AI could transform retail financial services by 2030 and beyond. It has been an honour. We drew on 140 written submissions and a wide range of expertise through meetings, panels and roundtables across financial services, technology, academia, consumer groups and regulation both in the UK and abroad. I am immensely grateful to everyone who contributed and gave up their time, and especially to the team and FCA colleagues who supported me.

One conversation stayed with me, from my mother: “Son, that AI is not going away now, you cannot switch it off, can you?”. As ever, her directness cut to it. It was also the view of nearly everyone we spoke to, pro-AI or sceptical. So, the question is no longer whether to allow AI, it is who is AI going to serve? My own view, as an almost former regulator, is that the answer must be citizens, and that empowering them is one of the UK’s great opportunities for economic growth. AI offers a once-in-a-generation chance to close the information asymmetries and frictions that have long left people making poor financial decisions. The right spur, in retail financial services, is ensuring consumers can make healthy ones; regulation, whether supportive or restrictive, should serve that outcome.

The pace has been the surprise of this Review. We began with a 2030 question, but the ground shifted as we worked: more than 20 frontier models have been released since I started in late 2025, and hundreds of smaller variants. AI is developing at a rate earlier technological shifts did not match. Firms are moving from systems that recommend actions to systems empowered and trained to take them, and consumers will soon gain agents that act on their behalf. That does not mean the FCA or anyone else should rush to write rules. It does mean the foundational work to support trusted, safe, and responsible adoption needs to accelerate, and I believe the FCA’s outcomes-based approach, which I was proud to help build through the Consumer Duty, provides the right basis for that work.

As autonomy grows, the nature of regulatory risk changes. As AI moves from recommending to acting, and firms and consumers delegate more, risks shift from harm within a single firm towards system-wide harms. One CEO quipped that we might need a Turing test for financial services. Practically speaking, we don’t know how far this goes and full autonomy seems unlikely to sit well with the current regulatory framework or UK societal appetites.

But we can prepare, and I believe an Agentic Supervisory Model giving the FCA the AI-enabled tools to monitor and detect emerging system-wide risks is a necessity. Striking the balance between enabling delegation and managing autonomy is the central challenge I believe the FCA is now committed to addressing.

What emerges is set out in the executive summary, which includes our recommendations. Taken as a whole, I believe the Review offers a practical agenda for the responsible development of AI in retail financial services which enables our world-leading industry to continue to grow and innovate, and meet the needs of all consumers.

Decisions on taking them forward rest with the FCA Board and Executive, and I wish them and the FCA every success in this endeavour. Go boldly where...

I was assisted by William Penwarden, Adam Boulton, Juan Lopez Pinzon, Laura Brown, George Cave, Oliver Parkes, Lucinda Craig, David Edward, Prasoon Chaturvedi and Arunita Roy.

I would like to thank George Langston and Chinmay Parikh for leading the team supporting the Review, and the wider team for their hard work, judgment and persistence throughout the project.

I would also like to thank colleagues across the FCA who provided technical advice, challenge and support during the Review.

Sheldon Mills

Chair,

Review into the long-term impact of AI on retail financial services

Executive summary

AI will transform retail financial services by 2030 and beyond. The strategic choices the FCA makes now will help shape that transformation.

The central shift is from human-led, episodic financial activity towards services that are AI-enabled, continuous and delegated. AI will operate inside firms, through consumer interfaces, across markets and within regulators. It will affect how products are designed, distributed, monitored and governed.

For consumers, AI could support better financial decisions and more personalised services. Unlike online search, it can combine a consumer's own information and data with wider knowledge to provide insights, intelligence and recommendations. Over time, these tools will evolve into AI agents that can execute tasks on consumers' behalf and help them manage their finances. Regulation can speed up or restrict that journey, and regulators now need to consider its opportunities, risks and benefits proactively.

AI can address longstanding weaknesses in retail financial markets, including the advice gap (only 9% of consumers use traditional advice), the protection gap (just 30% hold life or income protection), low switching, financial exclusion (around 900,000 are unbanked) and suboptimal saving (£300bn sitting in low-interest accounts). It can help consumers make better decisions, access more suitable products and manage their finances more effectively, bringing both personal and economy-wide benefits, since financial security, capability and access to capital are the bedrock of a modern economy.

Our [consumer research](#) found that trust, control and access are what must be right before consumers adopt AI finance agents. 1 in 5 UK adults are already open to AI making decisions for them, with demand strongest where choices feel complex or high-stakes, particularly debt advice, pensions and investments. Around 26% trust general-purpose tools such as ChatGPT, Claude or Gemini for financial advice, despite limited awareness that formal routes to recourse will not apply.

Firms, meanwhile, are adopting AI to increase productivity, reduce costs and improve operations. It can also support product innovation, ease barriers to entry for digital-native firms and reshape competition. Realising these benefits would strengthen the UK's position as a leading, trusted global financial centre.

AI capability is the systemic driver

Financial services is an information-intensive sector, and AI capability is advancing rapidly, transforming how firms and consumers use that information to make decisions. The frontier of AI is moving across many dimensions, with models now able to handle longer inputs, process more types of data, and reason through problems for better outputs. That trajectory of improvement looks set to continue.

Agentic AI has made it possible for models to use those tools to take action. Largely absent in previous years, it is now being piloted and deployed, allowing AI to take on increasingly complex tasks within firm and consumer workflows. A shift towards greater delegation is emerging, tempered by the constraints of model performance and the continuing need for human oversight.

Looking further ahead, future architectures are likely to help overcome the limitations of current models, making AI more efficient, explainable and adaptable, and potentially more general. On a longer horizon, AGI and quantum remain disruptive uncertainties. Their timing is unknown, but their consequences would be far-reaching.

Autonomy reframes risk and accountability

By 2030, AI will take on a greater share of the work in financial services, and this will be seen most clearly in how the human role changes. We set this out as an autonomy spectrum, along which humans move from performing tasks themselves towards setting boundaries, granting permissions and overseeing outcomes produced by AI agents acting on their behalf.

The spectrum runs through five roles. As an operator, the human uses AI as a tool. As a collaborator, the human and AI plan and act together. As a consultant, AI compares options and recommends, while the human decides. As an approver, AI prepares actions that the human authorises. As an observer, AI acts continuously within agreed limits while the human monitors outcomes.

Not every use case will move along the full length of this spectrum, and few parts of financial services will become entirely autonomous. The value of the spectrum is in showing how benefits and risks evolve as AI takes on more. Early questions of accuracy and reliance give way to more challenging questions about consent, accountability and redress.

The further a system sits along the spectrum, the more the human role shifts from carrying out individual tasks to setting the conditions within which AI systems operate. Accountability becomes harder to trace, especially where models change over time and firms depend more heavily on model providers and external AI agents.

Four shifts define the coming market

We see four systemic shifts that will reshape financial services to 2030.

1. AI will transform firms

Retail financial services firms have begun deploying AI for specific tasks, mostly keeping humans in the loop. This is changing. Firms are already piloting and rolling out use cases with greater autonomy, and by 2030 many could have moved significantly further along the autonomy spectrum, embedding AI into almost every function from customer support and underwriting to compliance, claims and product design. For leading firms, AI may become the main method by which they process information, serve customers, and evidence outcomes.

As autonomy grows, the role of people within firms changes, from operators close to each decision, towards collaborators, approvers and, eventually, observers who monitor outcomes and step in when systems move outside agreed parameters. This is a substantial organisational shift, requiring new skills and a clearer account of what human oversight actually involves.

Firm governance will extend existing model risk management to cover more complex systems and deeper reliance on third-party providers. Successful AI deployment should lift productivity and support economic growth, though the benefits will reach consumers only where firms remain accountable and markets stay competitive enough to pass them on.

2. Consumer journeys become agent-led

AI will reshape consumer financial journeys, with people increasingly delegating to AI applications that act on their behalf. Consumer demand is already emerging, suggesting a shift to agent-led journeys is credible.

Over time, AI systems will move beyond offering information and recommendations towards trusted AI agents that can act continuously for consumers within agreed limits, providing ongoing financial management and optimising people's financial lives.

If done well, this could help consumers achieve more while doing less, addressing long-standing problems such as low switching, advice and protection gaps, and improving outcomes for people with lower financial capability. There are both benefits and risks – for example, hyper-personalisation could help better match products to needs, but also enable bias, opaque pricing and personalised manipulation.

Access to useful and reliable AI services, trust and control will ultimately define consumer uptake and outcomes. Consumers will need to be able to oversee, understand and challenge AI-driven decisions, especially when things go wrong. Unequal access to high-quality applications risks widening inclusion gaps - but well-designed AI systems also present an opportunity to radically improve outcomes for those who need more support.

3. AI reshapes market power and competition

Effective competition can benefit consumers through lower costs, innovation and better products. But where it is impeded, or where there is significant market power or control of key inputs, it can be distorted, raising costs and lowering quality. AI has the potential to drive greater beneficial competition in financial services, yet it is also likely to create significant dependencies on suppliers who hold market power over firms' core operational capability.

AI could lower barriers to entry, enable new distribution channels and allow digital-native firms to scale rapidly. AI may also provide new ways for incumbents to leverage data, trust and bargaining power with AI suppliers. Control of the AI-mediated customer interface may become a major source of market power. As consumers rely on agents to search, compare and transact, the owner of that AI layer may influence which products are visible, how choices are ranked and where value is captured, shifting the customer relationship away from financial services providers.

This control could take several forms, from operating-system assistants and general-purpose apps to aggregators and consumers' own agents, each creating different competitive dynamics. Much will also depend on upstream technology markets, where competition among model providers, big tech and hyperscalers governs access to frontier capability, compute and data. Cost, concentration, vendor lock-in and sovereignty could all become critical to entry, innovation and resilience.

For regulators, this blurs the perimeter. Under an activity-based approach, general-purpose tools could shape financial decisions and competition without clear oversight. Whether this influence falls inside or outside the perimeter will shape how competition and consumer protection work in practice.

4. Threats and defences both accelerate

The same capabilities that promise to help consumers will also serve those who seek to target and defraud them. By 2030, AI is likely to amplify fraud and cyber risks, making attacks faster, cheaper, more scalable and more persuasive – and at the same time harder to spot and stop. Deepfakes, synthetic identities and personalised social engineering are taking fraud and cyber risks into a new era and changing how fraud and cyber-attacks are conducted. Existing weaknesses can be exploited far more quickly than before, and defenders will need to keep pace.

These risks extend across firms and platforms to payment rails, identity systems and technology providers, and across borders. As consumers increasingly rely on AI agents to manage financial decisions and transactions, risk will spread more quickly across an increasingly interconnected model-based system.

The same technologies used to attack the system can also be used to protect it. Defensive, supervisory and enforcement capability must evolve at least as quickly as the threat. To remain effective, firms, regulators and their partners will need access to many of the same AI capabilities as those used by attackers. They will also need to share the right information with those best placed to act, when it matters and before harm escalates.

The framework will need to evolve

We find that the overall regulatory framework remains sound. Its principles and outcomes-based approach, including the Consumer Duty, Senior Managers Regime (SMR), operational resilience and other key features, was designed to flex across changing business models. Engagement respondents did not seek changes to this system. They wanted clarity on how to interpret and govern increasing use of AI within the existing regime. We expect the FCA can provide this regularly through its supervisory, enforcement, authorisation, and its groundbreaking Scale-up Unit and AI Lab.

At firm level, AI creates new or amplified risks around governance, outcome testing and accountability, especially where models change quickly, decisions rely on complex supply chains, and firms automate key capabilities. No firm argued that the SMR accountability model should change, but many are considering how they continue to comply as automation increases. We agree that the SMR still applies, but firms and regulators will need to work through how it operates as AI systems become more autonomous. This links to our third priority recommendation.

The most significant change for the FCA to monitor is at system level. We usually understand systemic risk through a prudential lens, focused on financial stability and confidence. AI creates a broader challenge. Shared reliance on similar models, datasets and infrastructure providers could generate correlated behaviour, herding, opacity and common points of failure across the financial services system. A change or failure in a widely used model or service, or interactions between AI agents across firms, could be transmitted rapidly through firms and markets.

We also consider that, if trust frameworks were adopted and enough consumers delegated financial activity to AI agents, continuous movement of money could occur as those agents optimised for individual customers. This could, in theory, change pricing, products or services in markets such as savings and insurance. At market structure level, dependence on a few upstream technology providers could affect competition, resilience, sovereignty and smaller firms' access to advanced capability on fair terms.

Consumer and data protection risks will also grow. AI could widen access and improve personalisation, but it could also amplify bias, enable opaque or individualised pricing, exploit vulnerabilities and deepen exclusion where access to high-quality tools is unequal. That could diminish the opportunity to improve financial capability and decision-making.

Supervision must therefore extend beyond individual regulated firms to the financial ecosystem as a whole, with firm-level supervision supplemented by system-wide monitoring, coordination and oversight of shared models, infrastructure and cross-market dependencies. Regulators will need AI-enabled supervisory capability to identify cross-firm patterns, emerging harms and system-wide risks that no individual firm can see. Together, we call this the Agentic Supervisory Model and this forms priority recommendation 6, an AI-enabled approach that supports and enhances human supervision, so the opportunity from AI is not outweighed by risks of harm.

These four shifts: AI becoming core to firms, transforming consumer journeys, reshaping competition, and amplifying fraud risks, create both opportunities and risks and our recommendations provide a framework for the FCA to consider acting on them.

The FCA Board and Exec's opportunity

The FCA has a real opportunity to support the trusted adoption of AI, for the continued success of the UK's leading financial services sector, to enable UK consumers and small businesses to take healthy financial decisions, and to enhance the contribution of financial services to the UK economy.

The FCA cannot determine the speed of AI development or every market outcome, and some predict AI may never match human capability or reach the scale needed for the journey along the autonomy spectrum we outline. But all of our engagement, in the UK and beyond, showed that even the most sceptical believe AI cannot now be 'switched off'. It is here to stay.

So, the FCA should continue the excellent work it has done in this field already and work with Government, the Bank of England and PRA, and other sectoral regulators such as the Information Commissioner's Office (ICO), Competition and Market Authority (CMA) and Ofcom, alongside industry and consumer bodies, to influence how AI is adopted and how opportunity and risk are shared across financial services markets.








Striking this balance, between enabling innovation and greater delegation, while managing the risks of increasing autonomy, is the foundational work the FCA can commit to now.

It means putting the foundations of the future AI-enabled financial system in place over the coming years. Our seven priority recommendations work together as a system to adapt the regulatory framework, enable effective supervision, and support better consumer outcomes.

They are to:

- **Secure and adapt the regulatory perimeter**
- **Strengthen system-wide coordination and oversight**
- **Monitor the transition to autonomous models and adapt regulatory frameworks**
- **Scale up the FCA's AI Lab to support AI model and system innovation in financial services**
- **Enable the foundations for agentic finance**
- **Build and adopt an AI-enabled agentic supervisory model**
- **Develop a trusted public-interest AI-enabled financial capability service**

Figure 1.
How the recommendations operate together across an AI-enabled financial system

<p>Regulatory frameworks and perimeter</p>	<p> Secure and adapt the regulatory perimeter</p> <hr/> <p> Monitor the transition to autonomous models and adapt regulatory frameworks</p>
<p>Supervision and coordination</p>	<p> Strengthen system-wide coordination and oversight</p> <hr/> <p> Build and adopt an AI-enabled agentic supervisory model</p>
<p>Foundations and capability</p>	<p> Scale up the FCA's AI Lab to support AI models and system innovation in financial services</p> <hr/> <p> Enable the foundations for agentic finance</p>
<p>Consumer access and outcomes</p>	<p> Develop a trusted public-interest AI-enabled financial capability service</p>

The recommendations are mutually reinforcing

Together they create the conditions for safe innovation, effective oversight and better outcomes in an AI enabled financial system.

The AI autonomy spectrum

AI will transform financial services markets by 2030 and beyond. This report uses a spectrum of autonomy framework to explain that transition.¹ The spectrum should be understood through the changing role of the human: as AI systems move along this spectrum, people's roles change, from performing and directing individual tasks towards setting boundaries, granting permissions and overseeing and observing outcomes.

Human as **operator**, AI is mainly a tool that helps them understand, analyse or complete a defined task.

Human as **collaborator**, the human and AI plan and act together. When the human is a **consultant**, AI plays an active role by comparing options, recommending actions or preparing decisions, while the human gives guidance.

When the human becomes an **approver**, the AI prepares or initiates actions, but the human authorises key steps. Control shifts towards setting permissions and boundaries. Finally, when the human moves towards an **observer** role, the system acts continuously within boundaries set in advance, with the human monitoring outcomes rather than deciding.

The spectrum is not a prediction that all financial services will become fully autonomous, nor that every use case will move through each stage of the autonomy spectrum. It is a helpful way of distinguishing how both benefits and risks evolve as AI takes on a greater role, and how the human role changes over time.

When the human acts as operator, AI can improve efficiency, understanding and consistency, but raises risks around accuracy, evidence and appropriate reliance. As the human moves into a collaborator and then consultant role, AI can support better decisions and more tailored outcomes, while also introducing advice-like influence and risks of over-reliance or weak challenge. When the human becomes the approver, AI can reduce friction by preparing or executing actions with permission, but this shifts reliance onto how well permissions and safeguards are set in advance.

Finally, as the human moves towards an observer role, AI can deliver continuous, scalable support, but it also creates deeper risks around consent, accountability, auditability and redress, because the human is no longer making each decision, but setting parameters and overseeing outcomes afterwards.

The sections that follow use this distinction to explain why the same technology creates different benefits and risks depending on where the human sits in the process.

The spectrum of autonomy describes how the role of AI may evolve. The next section explains why we are likely to see potentially rapid movement along this spectrum and how it could reshape the wider retail financial services market. The principal systemic driver is the continuing advance in AI capability, supported by development in computing power, data and digital infrastructure.

¹ Feng et al, [Levels of Autonomy for AI Agents](#), July 2025.

Figure 2.



The systemic driver: Advances in AI capability

Key Findings

- AI capability is advancing rapidly and is likely to keep doing so to 2030 and beyond, transforming what firms and consumers can do with AI.
- Agentic AI is enabling a shift from assistance to delegation, with systems taking on longer tasks and more actions within firm and consumer workflows.
- Capable models still require controls for reliability, consistency, explainability and accountability, alongside human oversight as more is delegated to them.
- Future architectures are likely to change what is possible, making AI more efficient, explainable, adaptable and aware, and potentially more general.
- AGI and quantum are disruptive uncertainties with unknown timing but far-reaching potential consequences, so firms and regulators should prepare now.

AI capability continues to advance rapidly

Financial services is an information-intensive sector. Retail firms support consumers' everyday financial decisions, from accessing credit and insurance to managing money, and make judgements on affordability, risk, claims, and appropriate products. The quality of those judgements depends on the information available, the systems used to process it, and the people responsible for decisions. AI is reshaping how those decisions get made.

The current wave of AI differs from many forms used by the financial services industry in the past. Rules-based systems and task-specific machine learning models have long supported fraud detection, credit scoring, underwriting and operational decision-making. Many were designed for well-defined tasks and, once deployed, behaved in predictable ways.

Foundation models are less narrowly defined. They are general-purpose models trained on large and varied datasets and can be adapted to a wide range of tasks, including some they were not designed to perform. Large language models (LLM) are currently the most visible form of this technology, but the shift has extended beyond language. Many models now also process code, images, audio and video.

AI is moving deeper into financial services. A 2026 survey found that 81% of firms were adopting it at some level, with 40% at more advanced stages of scaling or transformation.² The most established use cases are still internal and operational, such as process automation, software engineering, data management and knowledge work. Customer support is the most common front-office use case, and fraud detection and credit risk modelling remain prominent in risk and compliance. Agentic AI, effectively absent from related surveys two years ago, is now being piloted or deployed by more than half industry respondents, although though most deployment remains limited in scope and subject to human oversight.³

² CCAF, [The 2026 Global AI in Financial Services Report – Adoption, impact and risks](#), 2026.

³ Bank of England and FCA, [Artificial intelligence in UK financial services](#), November 2024.

At the same time, model capability is improving.⁴ Current systems can handle longer inputs, work across different types of data, generate and test code, use external tools, and support more complex analysis than earlier models. Reasoning models, a category that barely existed before late 2024, can be given more time to work through a problem during use, improving performance on some planning and multi-step tasks.⁵ These gains are uneven, and performance varies by task, context and data quality. AI can now provide useful support across a wider range of financial services activity, although the need for controls remains.

The economics of use are changing as well. The cost of using models fell substantially in the two years following the release of ChatGPT, and continued to drop, making AI easier to deploy at scale.^{6 7} At the same time training the most advanced models is becoming more expensive. Overall spending appears to be rising as firms embed AI more widely, apply it to more complex work, and add the monitoring and integration needed in regulated environments.⁸

By 2030, AI is likely to be embedded across more of the systems through which retail financial services are delivered and controlled. Its impact will depend on model performance, and on how firms connect models to data, workflows, people and governance. The shift is a change in the operating architecture of financial services. It affects how work is done, how decisions are supported, and how firms and regulators can demonstrate that outcomes remain fair, explainable and resilient.

Agentic AI is driving a shift from assistance to delegation

AI becomes more significant for financial services when it moves from supporting action to taking action. Many current uses still leave the next step with a person. A model may summarise a file, draft a response or prepare material for human review, but a human decides what to do with it. Agentic AI changes that pattern by connecting model capability to tools, workflows and permissions.⁹ The system can pursue a goal, take steps within defined boundaries and interact with other systems on behalf of a firm or consumer.

In financial services, that brings AI closer to execution. An agent might prepare an application, update a customer record or initiate a payment. Some actions may be low risk and reversible. Others may affect access to a product or the outcome a consumer receives. The regulatory significance depends on what the system is allowed to do, how much it decides for itself, and where human approval sits.

4 AI capability is not plateauing, but accelerating, and several models meet or exceed human baselines on PhD-level science questions, multimodal reasoning, and competition mathematics. See Stanford HAI, [The 2026 AI Index Report](#), 2026.

5 OpenAI announced its o1 reasoning model in late 2024. See OpenAI, [Learning to reason with LLMs](#), September 2024.

6 There was a 280-fold drop in the inference cost of systems at a GPT-3.5 level. See Stanford HAI, [The 2025 AI Index Report](#), 2025.

7 Epoch.AI, [Trends in Artificial Intelligence](#)

8 The Economist, [Companies are scrambling to curtail soaring AI costs](#), June 2026.

9 OECD, [The agentic AI landscape and its conceptual foundations](#), February 2026.

The term “agentic” is applied loosely. Some systems described as agentic, including older rules-based workflows, involve little real autonomy. They are increasingly labelled agentic regardless, stretching a term that is losing precise meaning.¹⁰ Genuine agentic systems are different. They decide for themselves how to approach a task, which tools to use and when to seek human input. Among these, autonomy then varies along a spectrum, as noted earlier. At one end, a human operator uses AI to prepare information. Further along, the human acts as a consultant, with the system carrying out tasks and seeking guidance when necessary. At the far end, the system plans and executes multi-step processes, with the human as observer.

Two developments are pushing agents further along this spectrum. The length of task an agent can complete unaided is growing quickly. One study of frontier systems found the duration of work they can finish on their own, measured against how long it takes a person, has roughly doubled every seven months, and faster still more recently.¹¹ At the same time, agents are increasingly combined into multi-agent systems, where several specialised agents divide a task, coordinate and check each other's work, taking on longer and more complex objectives than one agent could alone. Together these stretch how long an agent can act on its own before a person needs to take control.

Where a person does take control is a question of risk. Retail financial services are likely to use a risk-tiered model, with more automation for tasks that are routine and easy to reverse. Stronger controls are needed where an action is high value or material to a customer outcome. Firms will need to decide where the human sits. That could mean approval before action, supervision while a workflow runs, or review by exception. Each choice affects accountability and redress.

Consumer-facing agents add a further complication. A consumer may begin by asking an AI application for information about a product. Over time, the same application may compare options, recommend a course of action and complete steps within parameters set by the consumer. Consent may become more continuous and conditional, with the consumer setting boundaries in advance instead of approving each step separately. As an agent takes a more active role on a consumer's behalf, it may also draw on a wider picture of the consumer than firms have used before.

Looking ahead to 2030, some of that wider picture may come from the interfaces and devices consumers use in their daily lives. Voice-enabled assistants are likely to be the more immediate route, allowing consumers to express questions, preferences and intent in more natural ways. Wearables, including smart watches, rings and glasses, could add further data about consumers' activity, location, routine and physical state. This data has not yet been closely integrated into financial services, but it may help build richer representations of consumers, enabling personal agents to interact with firms on the consumer's behalf.

¹⁰ Thoughtworks, [The dangers of AI 'agentwashing'](#), July 2025.

¹¹ Research on task-completion time horizons notes that the length of software task frontier models can complete autonomously at ~50% reliability rose from approximately 4 minutes in March 2024 to 12 hours as of February 2026 with early signs it could reach week-long tasks. See METR, [Task-Completion Time Horizons of Frontier AI Models](#), May 2026.

Over time, these technologies could become part of the ambient infrastructure around a consumer, informing more personalised and predictive decisions.

The move from assistance to delegation is therefore also a move in accountability. More financial activity may depend on interactions between systems, rather than direct human instruction at each point. Firms will still be responsible for the services they provide, including where they use third-party models or agentic tools. Regulation will need to preserve clear lines of authority, consent and redress as AI systems move closer to action. These issues are examined in our regulatory implications and recommendations sections.

Capable models still need controls

More delegation makes the reliability of the underlying model more important. Foundation models are useful because they can be applied across a wide range of tasks and often produce outputs that are convincing enough to be used in operational settings. That does not remove the need for firm-level controls. Outputs can still be plausible without being correct, a tendency commonly called hallucination, and firms may not always be able to show how a result was produced or why it should be relied on. In financial services, that can have direct consequences. A model may give a customer incorrect information, misread a policy term or support a decision that the firm cannot properly evidence.

Firms need controls around both the model and the system in which it operates. Outputs may need to be anchored to reliable data. Performance should be tested before deployment and monitored during use. These controls are most important where AI supports consequential decisions, such as whether to extend credit, pay a claim or uphold a complaint. Strong performance in general is not the same as reliability in the cases that matter most for consumers.

Consistency is a further difficulty. Models may produce different answers to similar questions, particularly where prompts, context or settings vary. That may be acceptable in drafting or analysis. It is harder to justify where firms need consistent treatment, clear audit trails and explainable outcomes. In regulated activity, a firm may need to show why one customer was treated differently from another. A useful answer is not enough if the basis for it cannot be reconstructed.

Models can also become less reliable over time. Financial products, fraud patterns and consumer behaviour change. A model tested at deployment may perform differently as the environment around it changes, a pattern sometimes called model drift. Firms need monitoring that can detect deterioration, identify emerging failures and trigger review. For higher risk uses, assurance cannot be a one-off exercise.

Human oversight remains important, but it is not a simple safeguard. A human reviewer can add judgement and accountability. Oversight can also weaken if reviewers are overloaded, lack the right information or rely too readily on model outputs. Firms need to decide where human approval is required, what reviewers should see and how challenge should be recorded.

Those choices should reflect the risk of the activity and whether an action can be reversed.

Controls are part of the deployed system. Evaluation, monitoring, audit and escalation shape whether firms can use AI safely in regulated services. Emerging AI insurance products are also forming part of the wider assurance landscape.¹² As models become more capable, assurance shifts towards the whole system, including the model, the data it uses, the workflow around it and the people accountable for the outcome.

Future architectures may change what is possible

The models used in financial services by 2030 may be quite different from those used today. Progress may come from larger general-purpose models, more specialised models, or improved verifiability. For firms, a practical question is whether these developments make AI more reliable in regulated settings. These directions sit at different stages of maturity, but they are advancing together rather than one after another.

Specialised models are already in use. Some are adapted for defined financial services tasks, such as fraud detection or payments.^{13 14} Others are smaller models designed for narrower uses. These models often cost less to run and are easier to test against known outcomes. They may also give firms more control over data and assurance than a general-purpose model accessed through an external platform.

The need for efficiency is driving new kinds of model. Retail financial services handle large volumes of data across long histories and real-time processes. Most foundation models use the transformer architecture, which becomes costly on long sequences. A newer class, state space models, scales more efficiently and is being used where transformers are expensive to run.¹⁵ As these mature, they could make real-time monitoring, fraud screening and transaction analysis more tractable. But an efficient model can still produce poor outcomes if the data, workflow or controls around it are weak.

Explainability is driving new approaches too. Neuro-symbolic systems, which combine machine-learned models with rules or formal logic, may help firms show how an output was reached and whether relevant requirements were applied.¹⁶ That could be valuable where firms need to demonstrate compliance. These systems may be harder to build, and may sacrifice some flexibility, but they lead towards AI that is easier to evidence and govern.

12 techUK, [Understanding AI Insurance and the role it plays in AI Assurance](#), September 2025.

13 Stripe, Revolut, Visa and Mastercard have all built models. For an example see Ostroukhov et al, [PRAGMA: Revolut Foundation Model](#), April 2026.

14 Thoughtworks, [The power of transaction foundation models in more detail](#), June 2026.

15 Mamba and Jamba are examples. Mamba is a state space model. Jamba is a hybrid that combines state space and transformer layers. See Lieber et al, [Jamba: A Hybrid Transformer-Mamba Language Model](#), March 2024.

16 Belle et al, [The Future Is Neuro-Symbolic: Where Has It Been, and Where Is It Going?](#), March 2026.

Some models keep updating after they are deployed, learning from new data rather than only at scheduled retraining. This approach is known as continual learning.¹⁷ Fraud patterns, customer behaviour and market dynamics all shift faster than periodic model updates, and models that adapt this way could respond sooner. That adaptability comes at a cost. A system that changes after deployment is harder to test and audit, so firms would need stronger monitoring and a reliable record of how it behaved over time.

Some models try to capture how an environment works, so that a system can anticipate the effects of an action before taking it. These world models are still early, and most progress has been in areas such as robotics rather than finance.¹⁸ But an agent able to model the likely consequences of an action before acting is better placed to stay within safe boundaries, which bears directly on how far agentic systems can be trusted to act.

The same idea extends to the market as a whole. A model that could represent how the financial system behaves, not just a single firm's processes, might let firms and regulators test how it responds before changes are made, including how interacting agents could affect competition, consumer behaviour or fraud. That is further off, and more relevant to regulators than to most firms, but it could let regulators test how the system responds before problems emerge, rather than reacting to them after they have occurred.

By 2030, firms are likely to use more mixed AI architectures. General-purpose models may sit alongside specialised and adaptive models, with verification tools and controls layered around them. Over time, and if these architectures prove reliable enough, capabilities now added around a model may move inside it. A neuro-symbolic model that can show its own reasoning, for instance, could remove the need for separate tools to check its outputs. Regulatory needs will remain. Firms will still need audit trails, accountability and oversight, even if the technical architecture changes.

AGI and quantum are disruptive uncertainties

How far AI capability could ultimately go is the largest uncertainty that the financial services industry and regulators face. Definitions of Artificial General Intelligence (AGI) vary, but most describe systems able to perform a wide range of economically valuable cognitive tasks, match broad human capability or improve their own capabilities over time.¹⁹ This report does not assume AGI will arrive by 2030, and its timing and form remain uncertain. But capability need not reach that threshold to matter. AI capability may be advancing faster than firms, supervisors and policymakers can respond.

¹⁷ See Yang et al, *Recent Advances of Foundation Language Models-based Continual Learning: A Survey*, January 2025. A related but distinct commercial direction is also emerging. Ineffable Intelligence, a UK lab that raised a record seed round in early 2026, is building systems that learn from their own experience, with the aim of discovering new knowledge rather than simply updating models over time.

¹⁸ See Li F., *A Functional Taxonomy of World Models*, June 2026. A number of model providers have a form of world model, including World Labs' Marble, Google's Genie, and NVIDIA's Cosmos.

¹⁹ Morris et al, *Levels of AGI for Operationalizing Progress on the Path to AGI*, September 2025.

Figure 3.
Model architecture evolution: From scaling to integrated intelligence



More general AI capabilities could emerge through several routes. Progress may come from further scaling of current models, from new architectures, or from combining the approaches described above. Each of those approaches answers a different need, whether efficiency, explainability, adaptability or the ability to anticipate consequences, and a system that drew on several at once could be more capable than any single model today. Models that are efficient at scale, can show their reasoning, learn from experience and anticipate the effects of actions might amount to a more capable and more general form of AI.

Today's dominant architecture should not be assumed to define the limits of AI capability through to 2030.²⁰

A further uncertainty is whether AI accelerates AI development itself.²¹ Models are already being used to write code, support research and test outputs. If they materially improve the process of building better systems, progress could become faster and less linear. The strongest version is recursive self-improvement, where AI improves itself with little human input.²² Whether or not it is possible remains contested. But even a partial version could leave less time for firms, supervisors and policymakers to adapt.

Large capability gains would affect the analytical core of financial services, such as research, compliance and market monitoring. Although they could support better products and more effective supervision, they could also enable harm at greater speed and scale and deepen concerns about concentration if critical capability remains held by a small number of frontier model developers. The range of possible outcomes is wide enough that firms and regulators should be planning for it, even where the timing is uncertain.

Quantum computing is separate from AI, but, as it matures, will affect both financial-services innovation and the foundations of digital trust. Its opportunity lies in using quantum mechanics to process some problems in ways conventional computers cannot. Quantum methods could improve optimisation, simulation and modelling, including in areas such as portfolio optimisation, fraud detection and asset pricing.²³ Any advantage is likely to be gradual and application specific.

The quantum threat to cryptography poses the more serious long-term systemic risk. Public-key encryption underpins most of digital finance, and a sufficiently capable quantum computer, known as a cryptographically relevant quantum computer (CRQC), could break the schemes now in common use.

20 The OECD consider four scenarios for how AI capability could advance or stall by 2030, and notes that further progress depends on scaling compute and data, algorithmic efficiency, and breakthroughs in areas such as memory, continual learning and causal reasoning. See OECD, [Exploring possible AI trajectories through 2030](#), February 2026.

21 Experts disagree about whether AI-assisted research automation could dramatically accelerate AI progress in the coming decade. See Bengio et al, [International AI Safety Report 2026](#), February 2026.

22 "Recursive self-improvement" (RSI) has no settled definition. A survey notes that several common definitions, taken literally, need not imply the explosive take-off often associated with it. See Cunningham T., [Definitions of Recursive Self-Improvement](#), June 2026. The idea of RSI is discussed by frontier labs, and pursued as a commercial goal by Recursive Superintelligence, a London lab founded in 2025.

23 FCA, [Quantum computing applications in financial services](#), October 2025.

No such machine exists yet, but the exposure is already present, and data captured today can be decrypted later, a risk known as 'harvest now, decrypt later'.²⁴

Firms should not wait for such a machine to exist before taking important steps. They can identify where vulnerable cryptography is used, build in the ability to switch cryptographic methods, known as crypto-agility, and plan migration to post-quantum standards.²⁵

That process will be long and complex, particularly for firms with legacy systems, outsourced technology and data with long-term confidentiality requirements.

AGI and quantum are strategic uncertainties. No one can say when either will arrive, but if they do, the effects would be far-reaching. That is why firms and regulators should prepare now rather than wait.²⁶ Neither needs to drive the analysis in this report, but both should shape the approach it recommends. A regulatory approach designed for 2030 should remain adaptive, evidence-led and able to respond if capability or risk changes faster than expected.

Implications for the rest of the report

The analysis that follows does not depend on predicting which model, provider or architecture will dominate by 2030. It starts from a narrower proposition. AI is becoming capable enough to change how financial services are delivered and controlled.

That proposition has consequences across the rest of the report. Firms may reorganise work as more tasks are supported or delegated to AI systems. Consumers may experience financial services through AI-enabled search, advice-like support, applications and switching. Markets may change as agents, platforms and customer interfaces play a larger role in distribution. Fraud and operational risks may also change as similar technologies become available to both firms and bad actors.

One regulatory issue is accountability under changing conditions. AI may improve service quality, reduce costs and strengthen risk monitoring. It may also make decisions harder to trace, weaken human oversight or increase dependence on a small number of technology providers. As AI systems move closer to action, firms and regulators will need to preserve clear responsibility for outcomes.

The following sections examine what increasing capability and delegation could mean for firms, consumers, competition and financial crime and regulation.

²⁴ NIST, [What Is Post-Quantum Cryptography?](#), 2026.

²⁵ The NCSC proposed targeting active PQC migration by 2031 and completion by 2035. See NCSC [Timelines for migration to post-quantum cryptography](#), March 2025.

²⁶ GOS sets out five scenarios spanning a slowdown, continuation or rapid acceleration of AI progress to 2030, and is intended to help policymakers stress-test plans against a range of possible futures rather than a single forecast. See GOS, [AI Scenarios 2030: Helping policymakers plan for the future of AI](#), June 2026.

System shift 1

The transformation of firms

System shift 1: The transformation of firms

Key findings

- Retail financial services firms have begun to deploy AI for specific use cases, broadly keeping humans in the loop.
- Firms are piloting and starting to roll out more autonomous AI use cases. By 2030 firms could have moved significantly further on the autonomy spectrum across a range of functions.
- There will be challenges for firms as they increase the level of AI autonomy, including the future role of human oversight.
- AI governance and model risk management will become a critical capability.
- Successful AI deployment by firms should have a positive impact on productivity and economic growth.

Engagement paper responses

- Responses consistently predicted that, by 2030, AI could materially change how firms design, operate and support financial services. This view was common even among firms that had not themselves made significant investments in AI.
- Stakeholders also highlighted that AI could build on earlier digital transformation. Previous waves of technology have already automated some aspects of financial services, but AI could extend automation into workflows that have historically remained more manual, complex or judgement intensive.

Retail financial services firms have begun to deploy AI for specific use cases

As set out in the AI capability chapter, retail financial services firms are using AI today for a range of tasks. The use of AI has expanded from non-agentic AI, such as coding assistance tools, to agentic AI, such as the use of AI in customer service functions.

Through our engagement paper responses and meetings, firms have set out that they are deploying AI across many functions today, and even firms which have made relatively little use of AI so far expect widespread AI deployment across financial services by 2030. This firm engagement has been consistent with a recent survey of retail and wholesale financial services firms on AI deployment.²⁷ In the survey, most responding firms reported having piloted or deployed AI in the categories of software development, of customer support and servicing, as well as fraud detection.

We have also seen some firms explore the use of AI in increasingly general business functions, such as in legal services or human resources. In addition, some firms have told us that they have already introduced some forms of multi-agentic systems, operating across internal processes.

Firms choose to deploy AI systems for different reasons, including revenue generation, cost reductions, and risk management. As AI continues to improve, and firms learn how to embed AI effectively, firms will have an incentive to deploy additional, more complex systems.

This current deployment provides an indicator of what more complicated future systems could look like, with multi-agentic systems operating across processes like product design, risk, compliance and assurance. By 2030, leading firms in their respective retail markets are likely to have embedded AI into almost every aspect of their business. AI may no longer be treated as an optional extra, but the method by which firms process information, serve customers, monitor risk and evidence outcomes.

Firms are piloting and starting to roll out more autonomous AI use cases

Initial deployments of AI by retail financial services firms have mostly had limited autonomy, with humans still being operators or collaborators. There are examples where autonomy has gone further. For example, agentic coding systems can allow developers to act as consultants or approvers.

There is a distinction between autonomy and automation. Where AI has been applied to an already automated function, such as an automatic assessment of a loan application, the role of the human operator has not changed, and so within the scope of activities done by humans, integrating AI into automated processes has limited effect on AI autonomy.

Future AI use cases could be more autonomous. Current deployments of AI set the foundations for later expansions.

²⁷ <https://www.jbs.cam.ac.uk/wp-content/uploads/2026/05/ccaf-2026-04-28-global-ai-in-financial-services-report-2.pdf> Cambridge Centre for Alternative Finance (CCAF); The 2026 Global AI in Financial Services Report

To set out what this might look like in practice, the table below sets out how firm activities may change where firms move across the AI autonomy spectrum.

Autonomy level	L1 Operator Human uses AI as a tool	L2 Collaborator Human and AI work together	L3 Consultant AI recommends, human decides	L4 Approver AI prepares, human authorises	L5 Observer AI acts within boundaries, human monitors
Operations	Developer asks AI to draft or explain code on demand Analyst requests summaries or charts from internal data and policy documents	Engineer and AI pair-program features, iterating and testing together Finance team co-builds FP&A forecasts, jointly refining assumptions	AI proposes vulnerability remediation strategy; security team steers priorities AI drafts contract mark-ups and pricing options; specialists guide and select	AI writes, tests and stages code; engineer approves each release AI diagnoses incidents and selects fixes; operations lead authorises execution	AI maintains and refactors codebases within mandates; engineers monitor outcomes AI curates knowledge bases end-to-end, choosing its approach; humans oversee
Front office	Adviser uses AI to summarise fact-finds or draft suitability letters Agent receives AI-suggested responses during live customer calls	Customer and chatbot co-build a budget or savings plan step-by-step Analyst and AI iterate together on an investment thesis	AI compares products market-wide and recommends; adviser directs and signs off AI designs personalised guidance journeys; humans shape and approve them.	AI agent executes savings switches via open banking; customer authorises each step AI renegotiates products across providers; customer approves final terms	Money-management agent optimises balances and bill payments across providers; customer monitors AI resolves servicing end-to-end, choosing resolution paths; humans review QA and complaints
Risk and compliance	Analyst asks AI to summarise a KYC file or explain a rule On-demand explanation of anomalies flagged in transaction data	Investigator and AI jointly build AML case files, testing hypotheses Compliance and AI co-draft horizon scans of incoming regulation	AI recommends SME and specialist lending decisions; underwriters guide edge cases AI leads alert triage and proposes escalations; humans set risk appetite	AI investigates suspicious activity and drafts the SAR; analyst approves filing AI prepares regulatory returns and validation reports; SMF signs off	AI adapts detection strategies and pursues investigations; humans monitor for drift AI monitors customer outcomes and selects interventions; committee oversees results

AI-enabled systems can increasingly support or perform complaints handling, underwriting support, claims triage, compliance monitoring, product documentation, quality assurance, internal policy interpretation, and software development. This shift is driven by the rapid advances in AI capability outlined in the AI capability chapter. It will increasingly be able to do this autonomously; in some cases, the AI system may prepare information for human review.²⁸ In others, it may recommend a decision, execute a low-risk action or escalate an exception.

Customer support is a clear example.²⁹ Earlier digital customer support depended on structured scripts or relatively narrow chatbots. AI-enabled systems can interpret more varied customer language and product information to support more flexible interactions. Some firms are already using more advanced chatbots and exploring AI-enabled voice interactions, while relationship-based businesses are embedding AI into adviser and client-servicing workflows.

However, AI has another clear advantage in that it can change how processes fit together. Earlier automation often sat within individual workflows, such as customer onboarding, handling complaints or detecting fraud. AI systems can connect these together and draw on data from workflows across the firm, monitor activity continuously, identify emerging patterns, prepare actions, escalate exceptions and learn from outcomes.

In this way, AI could also change how firms design and improve products as feedback loops across operations, servicing, risk and compliance provide better insight into product or service quality and consumer outcomes.

AI could therefore increase the responsiveness of financial services to market or customer changes. Instead of relying only on periodic review, manual sampling or fixed workflow rules, firms could monitor activity continuously and respond quickly when patterns emerge. This could support earlier intervention where customers are experiencing difficulty, where fraud patterns emerge, or where product performance suggests emerging harm. This continual monitoring could also allow some firms to spot emerging gaps in the market, which could form part of how firms design and roll out new products.

This does not mean all such work becomes fully autonomous. The quality of available data and the reliability of the AI system will determine whether autonomous outcomes are accurate. The degree of autonomy will also depend on the risk of the activity, the reversibility of the action, and the strength of governance. Of course, adoption in firms will be bounded not just by what AI can do, but by what consumers are willing to accept (as covered further in the next chapter).

²⁸ Cambridge Centre for Alternative Finance (CCAF); The 2026 Global AI in Financial Services Report, in particular that a material part of the vendor market sells beyond pure decision support.

²⁹ The CCAF's 2026 global survey found that "the leading front office use case is AI-powered customer support (74%)"

Lower-risk, routine and reversible activity is more likely to be automated end-to-end. Higher-risk activity, such as complex complaints investigations, reviewing disputed transactions, or investigating cases involving potential fraud or financial loss, is more likely to require human approval, review or escalation.

Indeed, we have heard from some stakeholders that they are thinking through how to construct an 'AI-native' financial firm from scratch.

How AI could be deployed across specific retail markets

By 2030, the impact of AI-enabled automation will vary market-to-market. Some market-specific potential market impacts are set out in Figure 4. (next page).

These examples are not intended to predict a single path for each market or firm. The core point is that AI changes the economics and feasibility of automation across a wider range of firm activity.

There will be challenges for firms as they increase the level of AI autonomy

As AI becomes more autonomous, the role of people inside firms changes. The Bank of England and FCA's AI Consortium have discussed "AI edge use cases", where firms have moved further along the autonomy spectrum. They may offer significant benefits, as well as challenges for governance, including where human challenge, approval or escalation is required.³⁰

At the lower end of the autonomy spectrum, people remain operators. AI helps them understand information, complete tasks, draft outputs or identify issues. The human remains close to the decision and carries out the next step.

As systems become more capable, people increasingly become collaborators or consultants. AI may compare options, recommend actions, prepare decisions or identify exceptions. The human role becomes one of challenge, judgement and review rather than direct production of every output.

Further along the spectrum, people become approvers. AI systems may prepare or execute actions within defined boundaries, with humans approving higher-risk steps or exceptions. At the highest levels of delegation, people may become observers, monitoring outcomes and intervening when systems move outside agreed parameters.

This is a significant organisational shift. Firms will need people with different skills: people who can understand AI-enabled journeys, risk and compliance specialists who can challenge AI-enabled workflows, supervisors who can interpret system behaviour, and senior managers who can evidence reasonable steps in automated environments.³¹

³⁰ Bank of England, Artificial Intelligence Consortium minutes February 2026, <https://www.bankofengland.co.uk/minutes/2026/february/ai-consortium-minutes-9-february-2026>

³¹ Financial Services Skills Commission, Annual Skills Report 2026, March 2026, <https://financialservicesskills.org/wp-content/uploads/2026/03/Annual-Skills-Report-2026.pdf>. 75% of surveyed UK financial services firms reported AI affecting their skills demand in 2025, with rising demand for AI literacy, data analysis and ethical AI governance.

Figure 4.
Potential market impacts of AI automated financial services



Many firms responding to the CCAF survey described accessing and training AI talent as a major challenge.

Human oversight and accountability will remain critical and many respondents underscore its importance. But it will not be enough to say that a person remains 'in the loop.' Firms will need to be clear about what the person is expected to do, what information they receive, when they can intervene, how challenge is recorded and how escalation works.

Strong AI governance becomes a core capability

As AI becomes more embedded, its governance becomes increasingly important. FCA regulated firms vary in size, capacity and governance frameworks. Some dual-regulated with the PRA or firms of size manage this through model risk management frameworks; others have governance or risk and compliance frameworks.³² The models shaping consumer outcomes including Consumer Duty monitoring, fraud and financial crime, complaints handling and affordability and creditworthiness assessments, as well as firms we prudentially supervise, often sit within these types of governance frameworks.

AI places pressure on these frameworks because models update continuously rather than on fixed release cycles, they draw on third-party inputs that a regulated firm did not create or build, and, as set out in the systemic driver chapter, they are probabilistic in nature. The management of AI means extending governance or model risk management frameworks beyond validation at the point of deployment towards live monitoring, tracking drift, model degradation and outliers.

Firms deploying AI at scale will need to adapt how they oversee their organisations. Managing integrated AI-enabled systems will become a core capability and may no longer work through periodic reviews but operate continuously alongside the systems it is designed to control.

This shift is already visible. Firms piloting and scaling AI in areas such as customer support and fraud detection report moving towards more real-time monitoring of system behaviour, building specialist governance capacity and extending oversight across increasingly complex third-party supply chains.

At the same time, AI can also help improve firm governance across the three lines of defence. In the first line, AI may support or execute customer operations, product delivery and servicing. In the second line, it may support risk, compliance, monitoring and policy interpretation. In the third line, it may support internal audit, assurance and testing. Across all three, AI can increase speed and scale, but it also requires stronger evidence of control. Moreover, as firms deploy AI at scale, governance must address not only performance but also heightened financial crime and cyber risks due to AI (see the financial crime and cyber risk chapter).

³² For example, dual-regulated firms such as banks have implemented model risk management frameworks, to comply with PS6/23.

Governance is likely to become an enabler of capability. Firms that can demonstrate auditability, explainability where needed, robust testing, clear permissions, effective monitoring and escalation will be able to deploy AI more confidently. Firms that cannot do so may be slower to adopt or may expose consumers and markets to greater risk. As our next chapter outlines, trust is a key barrier to customer adoption of AI use cases in finance, so acquiring a reputation for trusted AI processes could win business.

Finally, to operate safely at scale, firms need clear permissions, effective monitoring and auditability, supported by robust escalation mechanisms. These are not additional safeguards, but the conditions that enable AI to be deployed in regulated environments. As systems move from support towards execution, governance becomes a key enabler of capability and competitive advantage.

Managing dependencies in an AI-enabled firm

As AI becomes more embedded, firms rely increasingly on third parties, including AI model providers, broader AI infrastructure and specialist vendors.

Managing relationships and procurement with AI model providers and adjacent vendors therefore becomes a central governance task. Firms must oversee pricing, access to models, data sovereignty and relationships with specialist providers that sit between model developers and financial institutions. These dependencies also shape competitive dynamics, as access to leading models, infrastructure and vendors influences firms' ability to innovate, scale and compete (explored further in the competition chapter).

These dependencies become more consequential as AI moves closer to execution. Some firms may also allow customer agents or third-party systems to interact directly with their infrastructure. This increases the importance of managing not only internal AI systems, but also the terms on which external systems can access data, initiate actions and trigger firm workflows. Firms remain responsible for outcomes even where systems rely on external models or infrastructure.

Over time, these dependencies may evolve beyond traditional vendor relationships to broader agentic finance infrastructure that enables AI systems to act across firms and markets (see the agentic finance infrastructure annex).

Managing these dependencies therefore will become a core part of firm management, requiring oversight of performance, resilience and control across increasingly complex ecosystems. Managing this reliance from a regulatory and system-level perspective will be considered more fully in our priority recommendation 2 on strengthening system-wide coordination and oversight.

Implications for firm capability, productivity and economic growth

These firm-level productivity gains have wider economic implications.³³

³³ For more discussion on the difficulty of estimating future productivity in financial services see <https://www.bankofengland.co.uk/-/media/boe/files/financial-stability-in-focus/2025/financial-stability-in-focus-artificial-intelligence-in-the-financial-system.pdf> which discusses Accenture's work in the same area. This includes estimates of productivity gains of up to 30% to the banking and insurance sectors, and to firms operating in capital markets.

If firms can operate more efficiently and at lower cost, this can improve the value of financial services, either through lower prices, better products or wider access. In areas where the cost to serve has historically limited provision, AI may make it more viable for firms to reach underserved or marginal customer segments, supporting broader financial inclusion.

AI could also support better allocation of capital in two different ways. First, by changing how banks, building societies and insurers invest. Second, by changing where customers put their savings. Firms that can process information more effectively, assess risk more accurately and monitor outcomes more continuously may be able to price, lend, insure and invest more efficiently. This could improve the value of financial services and support wider economic productivity.

There may also be implications for UK competitiveness. If UK firms can deploy AI safely and effectively, they may develop capabilities, products and operating models that can be exported or adopted internationally.

This may require UK firms having and maintaining access to leading models. A trusted regulatory environment could support this by giving firms confidence to innovate while maintaining consumer protection and market integrity.

However, the benefits to consumers are not guaranteed. Efficiency gains may be retained by firms rather than passed through to consumers. Poorly governed AI could produce unfair, opaque or inconsistent outcomes. Shared reliance on the same models or infrastructure could create concentration and resilience risks. Smaller firms may benefit from cheaper AI tools, but they may also face higher relative costs for governance, testing, assurance and vendor management.

The productivity opportunity therefore depends on trust. AI-enabled transformation will support growth only if firms can deploy systems that remain controlled, accountable and resilient, and if markets remain competitive enough for benefits to reach consumers.

Implications for the rest of the report

These firm-level changes create the conditions for wider system shifts.

As firms embed AI into workflows, this shapes how consumers experience financial services and the availability and quality of products and services on offer. At the same time, as AI-enabled systems connect to consumer interfaces or consumer agents, they begin to influence how products are distributed and how markets function.

As firms adopt AI at scale, the speed, volume and complexity of decisions will increasingly exceed traditional supervisory approaches. This reinforces the need for an AI-enabled supervisory model, discussed in the priority recommendations chapter.

System shift 2

New consumer journeys

System shift 2: New consumer journeys

Key findings

- **AI will reshape consumer finance journeys:** consumers will increasingly delegate to AI applications to act on their behalf.
- **Demand for automated financial decisions is already emerging:** 1 in 5 consumers are open to AI making decisions for them, suggesting the shift to agent-led journeys is credible.
- **AI could help consumers achieve more while doing less:** AI already helps consumers navigate financial decisions and agentic AI could address long-standing issues such as low switching, advice gaps and low financial capability by acting continuously for consumers.
- **Hyper-personalisation brings risks and benefits:** tailored products may better match needs but could enable bias, opaque pricing and highly personalised manipulation or deceptive design.
- **Access, trust and control will define outcomes and uptake:** consumers will need the ability to oversee, understand, and challenge AI-driven decisions, while unequal access to high-quality AI applications could widen inclusion gaps.

Engagement paper responses

- **A balanced but cautious view:** most respondents took a measured view of AI's impact on consumers. They saw clear potential for improvement but stressed that strong safeguards are needed to ensure people are protected and treated fairly.
- **Key risks dominate the debate:** several concerns came up repeatedly, with respondents noting that the rise of highly personalised products, while bringing benefits, could drive risks of bias and exclusion, growing threats from AI-enabled fraud, and a lack of transparency. Many also questioned how consumers will be able to challenge outcomes or seek redress when things go wrong.
- **A once-in-a-generation opportunity:** at the same time, some saw AI as a major opportunity to improve financial services. It could help tackle long-standing problems and deliver better outcomes for millions of consumers.
- **Expanding access and making choices easier:** AI could widen access to financial advice, debt support and planning tools, helping people better understand and manage their money. It could also make it easier to find suitable products and switch between them, improving value and convenience.
- **Risk of widening inequality:** without the right decisions, there is a risk that AI - and access to it - could reinforce existing inequalities and leave some consumers worse off, rather than helping to close such gaps.

How AI will reshape financial journeys

By 2030, many financial consumer journeys may start with AI. Instead of searching for a product, comparing options and then choosing a provider, consumers may begin with an AI agent that helps them understand their needs, compare the market and act on their behalf.

AI, including publicly available LLMs from big tech platforms which sit outside the regulatory perimeter, is already helping consumers do more with less effort by turning complex information into clearer, more tailored guidance.

Over time, this may shift from AI supporting consumers as operators - helping them understand options and decide what to do - to consumers acting as collaborator or consultant to AI, with AI shaping decisions through recommendations and comparisons.

Further along the autonomy spectrum, AI agents may carry out financial services activities on the consumer's behalf - such as switching products or managing debt, savings or investments within agreed limits - with the human acting as an approver or, in rare cases, only an observer. We are already seeing emerging use cases across this spectrum.

As findings from our survey - covered more below - show, there is emerging appetite for automation. We expect to see more automated agent-led journeys where consumers act as consultants and approvers of AI agents if firms and regulators can improve reliability, data access, verification and identity, liability and redress, consumer support, and consumer trust. We provide a detailed examination of the building blocks of agentic finance in the agentic finance infrastructure annex.

Our research shows that agentic capabilities could increasingly be available through mainstream consumer platforms (Google, Apple, Amazon) and this will support research, comparison, planning and certain payments. But regular delegation for financial decisions of consequence remains uncertain and presents an opportunity for the FCA - as we set out in our recommendation on enabling the foundations for agentic finance (priority recommendation 5) - to support the development of trusted agent protocols which would underpin use of agentic AI in the consultant and approver stages of our autonomy spectrum.

Consumer outcomes will depend not only on AI capability, but also on the interfaces through which AI is delivered. Bank-controlled applications, general-purpose platforms, operating-system assistants, aggregators and independent agents each create different incentives, levels of transparency and accountability, as discussed in the competition chapter.

In the meantime, we anticipate that individual consumers will increasingly create their own agents using AI systems and delegate to them aspects of their personal financial management.

We have seen examples of this in our research - for example, investment platforms Robinhood and Public allow customers in the US to connect their own agents and enable them to execute trading strategies based on set criteria.^{34 35} Robinhood also allows customers' agents to spend on their behalf via credit card, again enabling agents to take action based on customers' instructions. This is covered further under System Shift 3, where we discuss customers choosing independent AI agents.

What consumers are telling us

Of course, such journeys can only take place where there is consumer demand. With limited existing evidence of demand, particularly for more advanced, agent-led uses, the Mills Review commissioned new research from Yonder Consulting to fill this gap.

This nationally representative survey of 5,026 UK adults assessed current consumer attitudes and behaviour relating to AI in financial services, including delegating financial decisions to AI systems. This followed eight consumer focus groups, run by the FCA in London and Leeds, which informed the design of the survey and highlighted a range of attitudes. Survey findings, together with selected qualitative observations, are in the consumer survey and focus group annexes.

We found that AI use is widespread among UK adults, but financial services adoption is less common and mostly limited to getting information or suggestions. Financial services use cases are higher in areas such as debt advice, pensions and investments, and lower for more routine products. Demand seems strongest where decisions feel complex, high-stakes or hard to navigate.

In our qualitative research, people were often uncertain about specific benefits but spoke generally about *'a bit more financial freedom'*, and how *'it would help me be more organised, hopefully save money and spend more wisely.'*

Multiple general concerns about AI are for now holding back wider adoption. 68% of consumers expressed concern about data misuse, 67% about lack of protection and 65% about concentration of power among a small number of large organisations, while 45% saw no benefit of AI in financial services.

Currently, consumers remain most comfortable with AI as an aid, rather than a decision-maker. Practical benefits of convenience, where AI explains complex information or reduces effort are compelling, but many consumers are less comfortable where AI moves closer to advice, decision-making or control over money.

Despite this, some behaviours point towards more advanced use. Among consumers who use AI in financial services, 24% report uploading personal financial data and 13% say they would be willing to grant real-time access to their financial information. This willingness increases with AI familiarity: 31% among those who use AI daily for any purpose and 36% among those who already use AI for personal finance or financial products.

³⁴ Robinhood, "[Robinhood is now open to agents](#)", 2026.

³⁵ Public, "[Agents. For your portfolio](#)", 2026.


Figure 5.
Consumers tend to agree that AI in FS will save them time and money, but once autonomy increases so do concerns around scams, fraud and stress

Chart showing % agreeing with benefit minus % agreeing with drawback for each level.



C3. What impact would using this have on you and how you manage your finances?
 Base: all respondents (5026) L1 = Assistive L2 = Agentic L3 = Autonomous

Figure 6.
Headlines on usage, benefits and concerns

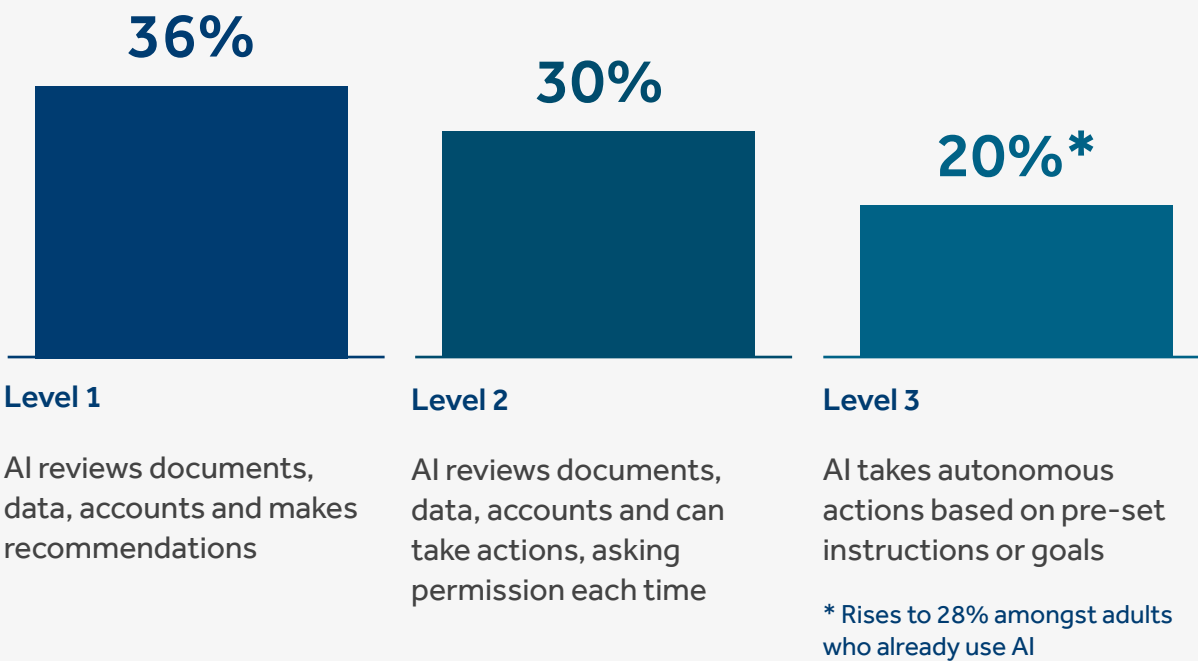


*16% of the 92% who did one or more financial task
Base: 5026 UK Adults, 3473 have used AI, 900 have used AI for personal or financial products.

Figure 7.

20% of consumers would use AI that acts autonomously within pre-set goals. Although reported willingness to use AI currently declines as autonomy increases, it is notable that one in five consumers would consider autonomous AI

Chart showing % saying they would definitely / probably use something like this, if it were available.



C2. How likely would you be to use something like this, if it were available to you now?
Base: all respondents (5026)

So, even at this early stage, a significant minority are moving towards more integrated and potentially higher risk uses that raise questions about data protection, consent and accountability.

Importantly, there is evidence of underlying demand for automation. When presented, with realistic scenarios involving delegation of financial decisions to AI systems, 20% of all adults, and 28% of existing AI users, responded positively to the most advanced concepts tested. Similarly in our focus groups, one member said, *'When I become more accustomed, I can open it up to do more things for me.'* The implication is that behaviour could shift quickly as capability improves and concerns are addressed.

*'I completely believe this is possible
and would use it if it was guaranteed to be secure.'*

*'I would find this really helpful. Sometimes finances can be really confusing
so if this can make it easier, then I would use it'*

Trust will be a defining factor in adoption

Consumers are open to the concept of delegation in personal finance even if they are not yet confident enough to adopt it at scale. Among AI users, 75% trust AI to explain things clearly and 71% to provide useful guidance. However, trust falls sharply when it comes to AI avoiding misleading information (23%) and handling personal information responsibly (21%).

Confidence in AI remains limited: around 28% of AI users strongly trust tools like ChatGPT to provide accurate information (trust 'a lot' or 'completely').

Understanding of risks is also weak: when these tools are used for financial advice, only around two in five people correctly identify the level of protection available, while a majority either misunderstand it or are unsure.

Protection and accountability are central to whether the shift to AI-enabled financial journeys improves outcomes or amplifies harm. People believe more safeguards should be enacted to protect them from harm. For example, if AI-enabled fraud is not effectively managed, trust in AI-mediated financial services could be significantly undermined (see the financial crime and cyber risk chapter).

*'If it does something and it cost me loads of money, am I protected
or is it, kind of, you signed this secret waiver?'*

*'If it's made a poor decision, you invest all your money. You lose it or whatever.
Is there regulation or any back up? Or is it on you still?'*

Figure 8.
Four current attitudinal cohorts



Aware of AI limitations - don't trust AI to admit when it doesn't know something and to avoid providing misleading information.
 Less aware of AI limitations - do trust AI to admit when it doesn't know something and/or to avoid providing misleading information.
 High AI confidence - report feeling very or quite confident in their ability to use AI or AI-enabled apps.
 Low AI confidence - report feeling not very or not at all confident, or say they don't know, when asked about their ability to use AI or AI-enabled apps.

Figure 8. (previous page) shows four attitudinal cohorts of consumers according to their confidence in AI and degree of critical engagement with outputs, based on our representative Yonder survey. A small group appear more willing to rely on AI outputs uncritically, particularly where they trust it to avoid misleading information. These consumers are more open to autonomous financial services applications but may be less attuned to the risks. This chart is important because it points to how uneven levels of trust limit the development of this market.

Less confident consumers may not adopt AI and miss its benefits, while more confident users may rely on it inappropriately and be exposed to harm. If this persists, overall demand may not develop at scale. This matters because, if adopted safely, agentic finance has the potential to improve outcomes for many consumers, as set out below.

Why agentic finance could improve outcomes

As highlighted in Figure 9. (next page), AI agents acting on behalf of consumers could help tackle some of the long-standing inertia that prevents consumers from getting better outcomes in financial services.

As noted across this review, switching remains low in several markets. A trusted agent could monitor products continuously, compare the market at renewal points, prompt action when a better option is available, and carry out routine tasks within limits set by the consumer.

The same logic applies to advice, saving and investing. Few UK adults receive financial advice each year or hold investments. AI could help consumers understand trade-offs, identify when guidance or advice may be needed, and help manage savings and investments more actively.

It could also help address protection gaps. While 84% of adults hold a general insurance policy, only 30% hold protection policies such as life or income protection.³⁶ A personal agent with a fuller view of a consumer's circumstances could help identify gaps, explain options, and prompt consumers to consider cover at relevant life moments.

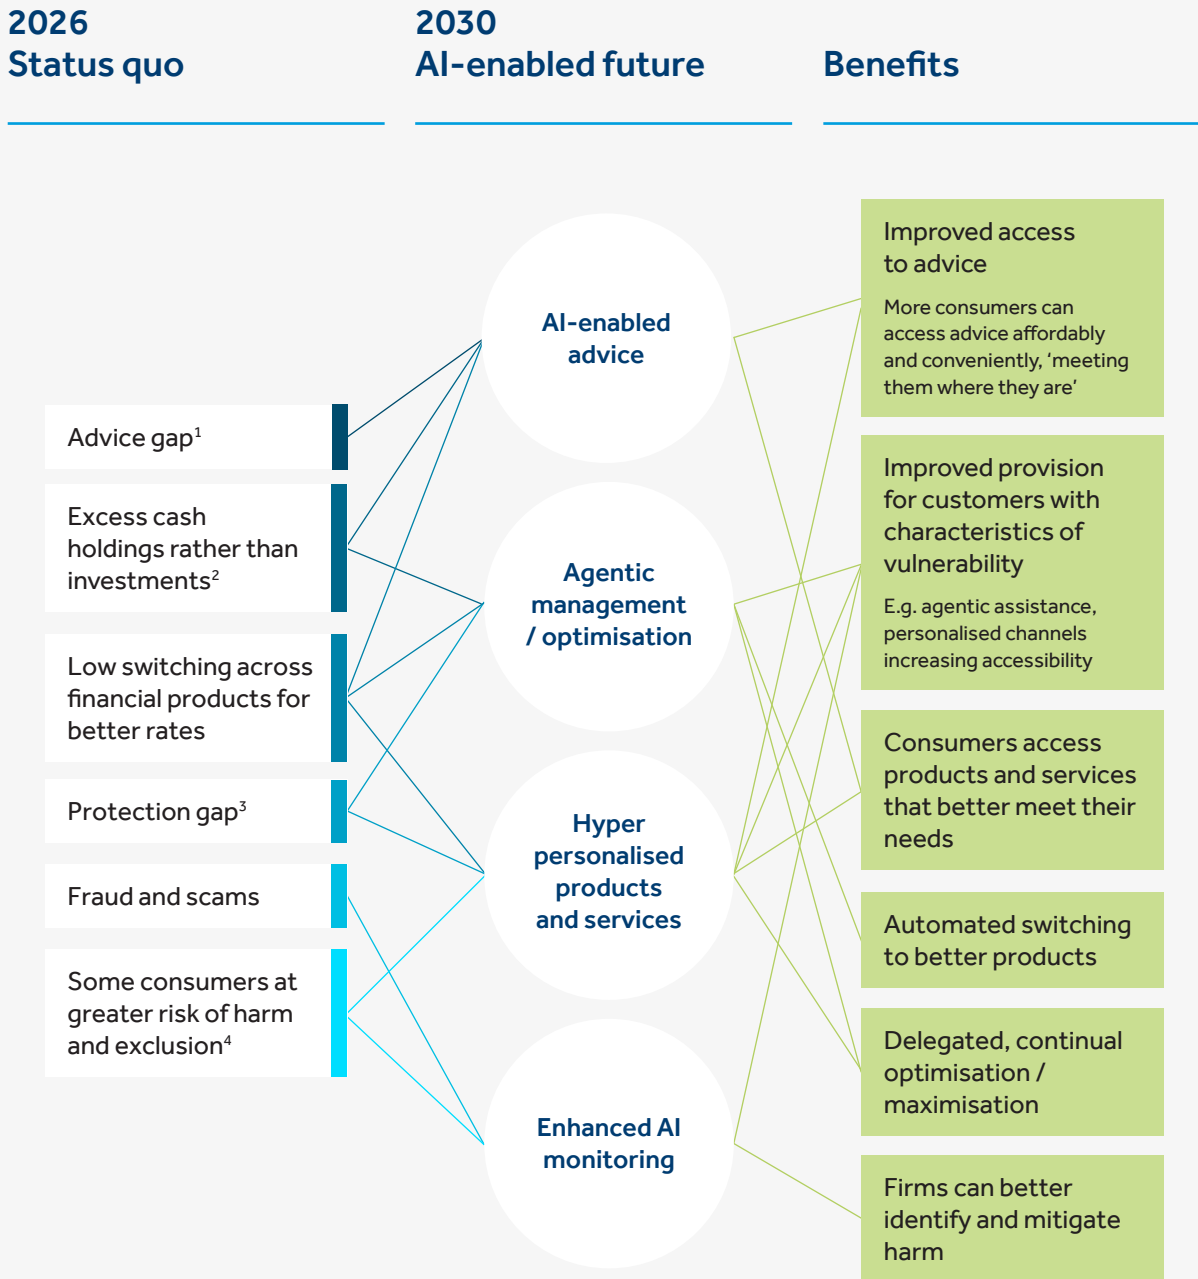
The most powerful consumer benefit is not simply convenience. AI agents could help deliver good outcomes that do not depend on consumers having higher capability and engaging constantly. If designed well, they could help those who are less financially capable, time-poor, or less confident to act more often, review decisions more regularly, and secure outcomes that currently require too much effort for such consumers.

However, this will change who shapes the consumer journey. Currently this is financial firms, but in the future agents could become powerful gatekeepers, influencing what consumers see, which products are compared and how choices are framed. We discuss this area more under the competition chapter.

The shift is likely to happen in stages. First, consumers act as operators, using free or low-cost tools to explain financial terms, summarise options or prepare questions.

³⁶ FCA, [Financial Lives 2024 survey - General insurance & protection: Selected findings](#), p.29.

Figure 9.
AI-enabled financial services: from today's gaps to a better 2030



1 Only 9% of people receive financial advice annually [FLS, p.39]
 2 Only 35% of adults hold investments and of those with £10,000+ in investible assets, 61% held all or at least three-quarters of these assets in cash [FLS, p.15]
 3 Only 28% of renters hold contents insurance [FLS, p.29]
 4 Customers with characteristics of vulnerability face barriers engaging with financial services - for example only 45% of those with low financial capability agree there are suitable ways to contact a financial services provider when they needed to (65% all adults) [FLS, p.106]

Then they use more capable tools, collaborating with AI to compare products, build budgets or organise debt repayments. Later, consumers may allow agents to monitor their finances, prompt action, or execute tasks such as switching accounts, reallocating savings or managing subscriptions. By this point, consumers may have delegated most ongoing action to AI, signing off actions at key points, or reviewing activity logs of what their AI has done for them based on agreed parameters and preferences.

In practical terms, a consumer might begin by asking an AI agent whether they are paying too much for insurance. The agent could explain the policy, compare alternatives, identify exclusions, and later switch the consumer to a better product within limits they have set.

The consumer may still be in control, but the work of understanding, comparing, and acting has shifted away from them. This could bring real benefits. It could reduce inertia, help consumers act on information, and make markets more competitive if AI services search across providers and highlight better options. But it could also weaken direct relationships between consumers and firms, reduce brand loyalty, and introduce new commercial incentives into the journey. So, these benefits depend heavily on how these agents are designed, governed, and embedded within markets. Competition dynamics, discussed in the competition chapter will play a big role in shaping this landscape for consumers.

How AI could change consumer capability

Beyond changing how consumers interact with markets, AI could also change how they think, decide and build financial capability over time. Consumers who struggle with complex financial decisions may be able to use AI to understand information, compare options and act with confidence. In journeys at the higher end of the autonomy spectrum, consumers may be able to achieve more while doing less themselves.

That could be empowering, especially for consumers in vulnerable circumstances or those with low literacy, numeracy or confidence. But it also creates a risk that some consumers become over-reliant on AI, trust poor outputs, or lose the ability to challenge decisions made on their behalf. Academic and consumer group stakeholders raised this as a novel risk - whether repeated use changes how consumers think, decide and act over time.

The positive side of this change is that AI helps people who find finance confusing to ask better questions, see clearer choices and take action they might otherwise avoid. At the further reaches of the autonomy spectrum consumers may not need to build deep financial capability if a reliable agent can help them achieve their objectives safely.

The risk is that this cognitive offloading changes behaviour.³⁷ Consumers' critical thinking may be affected, meaning they become less able to judge whether AI outputs are reliable, less likely or able to question recommendations, or less prepared to act independently if AI services fail or become unavailable.³⁸

³⁷ Kosmyna et al., "Your brain on chatgpt: Accumulation of cognitive debt when using an ai assistant for essay writing task." [arXiv preprint arXiv:2506.08872 \(2025\)](https://arxiv.org/abs/2506.08872)

³⁸ Shaw and Nave, "[Thinking - Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the rise of Cognitive Surrender](#)", 2026.

This is not a reason to reject AI, but it does mean that design choices should encourage appropriate scrutiny rather than passive acceptance. In the future, defining the most important elements of financial capability will, in part, be a response to the design and abilities of AI models and their incentive structures.

*'I don't want to become brain dead
because the AI is making all the decisions for me.'*

We also heard from the debt advice and consumer support sector about cases where consumers had first sought help from general-purpose AI services and received incorrect information. Advisers then faced the challenge of steering those consumers back to a safer course of action.

This reflects the continued influence effect where misinformation can continue to shape thinking and decisions even after it has been corrected. In an AI context, that risk may be amplified where consumers overestimate an AI application's accuracy, attribute abilities to it that it does not have, or treat a fluent answer as a reliable one.

We should therefore avoid treating capability, vulnerability and access as separate issues. By 2030, financial capability may be inseparable from digital capability, AI literacy and access to high-quality AI services. The policy question may become whether it is important to prioritise financial capability, or more important to prioritise consumers being able to use, question and control systems that are engaged to act in their interests.

To help examine these policy questions, the FCA could consider AI's impacts on vulnerability and capability in any future work, for example, in any review or clarification of relevant rules and guidance, such as the Vulnerability Guidance and Consumer Duty.^{39 40} The FCA could also make use of opportunities available to help firms understand how its expectations apply in relation to AI and how they act to deliver good outcomes for customers in vulnerable circumstances in AI-based journeys, for example by offering clarity through good and poor practice publications.

Three implications for consumers

The expected shift to AI-mediated journeys points to three core consumer issues: **who benefits, how decisions are shaped, and who is responsible when things go wrong.**

1. Who benefits and who is left behind

AI could make financial services easier to access and use. Personalised interfaces and agents could help consumers understand products, manage money, seek debt support or take action that previously felt too difficult.

Despite this, in our survey, although we saw positive signals about adoption, we also saw tension and resistance.

39 [FG21/1: Guidance for firms on the fair treatment of vulnerable customers](#)

40 [FG22/5: Final non-Handbook Guidance for firms on the Consumer Duty](#)

45% of consumers said they see no benefit in using AI for day-to-day financial management, while 24% said nothing would encourage them to use it. Consumers saying this were generally older and on lower income. Access to good AI may therefore become a new divide.

If the best services are paid for, difficult to use or poorly designed, or simply not taken up, some consumers who could benefit most may be least able to access them. Firms will need to design inclusively and avoid shutting consumers out of AI-enabled support where it could improve outcomes. There is also a threshold problem where consumers who could benefit the most might struggle taking the first step towards using AI safely without help.

Work at the FCA and within the industry may therefore need to cover financial capability, digital access and AI confidence together.

Access is not just about price. Some consumers may not want to use AI, not trust it, or need non-digital channels because of disability, low confidence, language barriers or vulnerability. If markets become AI-first by default, those consumers could face reduced choice or worse service unless firms design AI routes inclusively to meet the different needs of their customers and ensure access via other routes (i.e., not digital or AI-based) where needed.

There is also a risk of a two-tier market. Our survey showed that 11% of AI users paid for access to AI (outside of work) and we know that some AI providers charge more for better capability for some tasks. As is shown in the competition chapter, the cost to use AI might increase for firms and consumers, and this could further exacerbate the risk of a two-tier market where consumers with access to more capable AI systems or interfaces may receive better answers, better comparisons and better ongoing management of their finances. Consumers without those systems could be left with poorer information, more friction and less ability to benefit from competition. That would turn access to more capable AI services into a driver of financial inclusion.

2. How decisions are shaped

AI can personalise products, pricing and support in ways that may genuinely improve outcomes. Better matching could help consumers find more suitable products, switch more easily and receive services that better meet their needs.

But the same AI services could make markets more opaque. Consumers may see only the options selected by an AI system and not understand why a price, product or recommendation has been shown to them.⁴¹ AI's deeper analysis of much larger amounts of data could also amplify bias, exploitative pricing or exclusion unless firms can monitor, explain and correct outcomes.

⁴¹ Ada Lovelace Institute, "[The dilemmas of delegation](#)", 2025.

Focus group cynics said:

'With certain AIs, they're only using the pathways they're given. They're not using all of them or they're going to make more money using certain pathways.'

'If they're getting everyone's exact financial decisions and they know exactly what you want to do with your money and where to go, I think they could use it for their benefit.'

Hyper-personalisation is the clearest example of this. At its best, it could make products and services more relevant, encourage firms to innovate to serve new consumer niches, adapt support to individual circumstances, and help consumers find products that better meet their needs. At its worst, it could create a market where consumers see only a narrow, personalised slice of options and cannot tell whether they are receiving a fair deal. The impacts of hyper-personalisation may vary depending on the interface model. Firm-controlled applications may have stronger incentives to steer consumers towards their own products, while platform, operating-system or aggregator models may shape choice through ranking, defaults, commercial arrangements or technical integrations.

Bias and discrimination risks become harder to manage where firms use larger datasets, more complex models and variables that may act as proxies for protected characteristics or characteristics of vulnerability. Some consumers could face higher prices or reduced access because the model has identified patterns linked to disadvantage rather than legitimate risk. This risk was often raised by Engagement Paper respondents. Firms will need to be able to distinguish accurate risk-based pricing from unfair or discriminatory outcomes.

Design also matters. AI interfaces can make consumer journeys different for each person, which could improve relevance but reduce transparency. This makes auditability, explainability and clear consent points more important.

Indeed, one of the more distinctive risks in AI-mediated journeys is the potential for highly personalised forms of deceptive design. In today's digital markets, firms can already use presentation, defaults, friction and prompts to steer consumers towards particular choices. AI could amplify this risk by enabling firms to identify which forms of presentation, prompts or friction are most effective for each individual consumer. These steers can be adapted in real time and deployed at scale. This may allow firms to exploit consumers' immediate circumstances, behavioural data, preferences, or vulnerabilities in ways that are more precise and difficult to detect than in conventional digital consumer journeys.

This matters because a commercial AI interface might appear to provide neutral assistance while the system may be optimised for engagement, conversion or revenue.

It could frame options in ways that nudge a consumer towards a sale, discourage switching, downplay exclusions, or encourage a consumer to proceed where advice or additional support would be more appropriate.

Where a consumer is only an approver or observer for their AI agent the risk might be even greater - less direct interaction in a journey means consumers might be less likely to identify when they are receiving poor treatment. AI agents may also be susceptible to manipulation or steering by firm's systems or other agents. Conversely, an independent AI agent - for example, where a consumer creates their own agent rather than using one from a firm - could protect consumers by filtering manipulative prompts, managing and controlling access to data, comparing options consistently across the market, and counteracting behavioural biases.

This might offer even greater protection for consumers.

This makes testing, monitoring and auditability central to ensuring that AI-enabled design improves consumer outcomes and doesn't restrict consumer choice. It underscores the need for system wide supervisory approaches by the FCA as set out in our recommendation on building and adopting an Agentic Supervisory Model (priority recommendation 6).

3. Who is responsible when things go wrong

High levels of concerns about AI in financial services were recorded in the Yonder survey. Around three-fifths of respondents reported concern across most risks tested. Concern is particularly high in relation to misuse of personal and financial data (68%), lack of protection if something goes wrong (67%) as well as concentration of power among large firms (65%). Similar levels of concern are reported for fraud, biased recommendations and financial loss. Importantly, these concerns extend beyond technical performance. They centre on accountability, incentives and consumer protection. Adoption is constrained more by confidence and perceived risk than by technical capability. We therefore stress the importance of trusted agentic protocols in our recommendation on enabling the foundations of agentic finance (priority recommendation 5).

As mentioned, dynamic AI consumer journeys could improve support and make services easier to use. In some current journeys, these are enabled through partnerships between regulated firms and model providers – for instance where financial services firms have integrated apps into ChatGPT. These integrations could make it harder to see what happened when something goes wrong and consumers may not know whether responsibility sits with a financial services firm, a model provider, a platform, or an agent acting on their behalf.

Trust will depend on consumers being able to understand decisions, challenge outcomes and access redress. Firms will need auditable records, clear complaints routes and enough understanding of AI outputs to explain decisions and mitigate harm.

AI could also change the complaints system itself. Used well, it is already helping consumers improve the structure and clarity of their complaints, and in future it could help reveal problems earlier, improve access to the redress system for those less able to represent their interests, and help consumers to understand their rights and gather evidence.

Used poorly, it is already lowering the quality of some complaints with irrelevant legal arguments and unhelpful AI-generated content that puts extra strain on firms and the Financial Ombudsman Service. In the future, it could make complaints cheaper and easier to generate, potentially increasing volumes of low-quality, automated complaints that increase operational strain and make genuine harm harder to identify.

Contestability will require several things: consumers need to be able to recognise that something has gone wrong; firms need records that show what happened; the decision logic must be sufficiently understood to explain the outcome; and consumers need a clear route to complain, including where the journey involved parties outside the traditional financial services relationship. We discuss this in more detail in our recommendation on enabling the foundations for agentic finance.

Looking ahead, some complaints may become agent-to-agent interactions, with consumer agents raising issues and firm agents triaging or resolving them. This could mean issues get resolved agentially without the need for them to become a full complaint. It could also improve efficiency, but only if systems are reliable, interoperable, auditable and capable of escalating complex or sensitive cases to people.

The FCA should consider the impacts of these changes on the redress system, incorporating them into its work with HM Treasury, Financial Ombudsman Service, Financial Services Compensation Scheme, and firms. It should pay particularly close attention to how agentic AI might affect rapid scaling of complaints issues, including by professional representatives.

Implications for the rest of the report

AI-enabled consumer journeys bring the wider system shifts together. As consumers use AI to understand, compare and act, firms will need to adapt how they design and distribute products and services, and how they evidence AI-driven journeys and decisions.

These journeys will also reshape competition. If AI interfaces determine what consumers see, compare and the actions they take, market power may shift towards the platforms, agents and applications that mediate consumer access. This is the central issue in the competition chapter.

They may also increase consumers' exposure to fraud and manipulation, cyber risks and failures of accountability, reinforcing the need for the coordination discussed in the financial crime and cyber risk chapter.

Overall, this chapter shows why the regulatory response should focus on trusted adoption - enabling AI to improve access, capability and outcomes while preserving consumer control, inclusion, contestability and redress.

This underpins the later recommendations on the perimeter, agentic finance infrastructure, AI-enabled supervision and inclusive access to trusted AI support.

System shift 3

A reshaped competition landscape

System shift 3: A reshaped competition landscape

Key findings

- As AI becomes an increasingly important input into retail financial services, competition within the AI supply chain will affect downstream competition between firms and consumer outcomes.
- AI facilitates additional interfaces for consumers to buy and access retail financial services. Competition will occur within and between these interfaces.
- Control of and access to AI-mediated consumer interfaces will become a major source of market power and reshape competitive dynamics.
- AI may shift the nature of intermediation towards integrating with agents, platforms, and facilitating data-access between financial services and AI.
- Firm-level AI capability will become a core competitive advantage: AI may lower barriers to entry and growth for new and smaller firms, while the impact of scale on data, trust and vendor access may favour incumbents.

Engagement paper responses

- Respondents agreed that whoever controls the customer interface will likely have significant market power.
- Firms across different markets think that there will be significant effects on intermediary markets, with disruption to brokers.
- The majority of regulated firm respondents raised concerns around a level-playing field due to AI systems being able to operate as tailored search, information and recommendation tools outside the perimeter
- AI-native firms are expected to be able to benefit from lower operational costs but may struggle to get market share.
- Responses highlighted that there was concentration throughout the AI supply chain, with limited numbers of hyperscalers, model providers, and chip manufacturers.
- The fixed costs of some financial investments in AI meant that AI might systematically benefit large incumbents.

As we set out in the firms chapter, retail financial services firms will make increasing use of AI. This chapter explores three competition channels which will affect consumers.

Firstly, AI model providers face multiple options for monetising the use of AI within retail financial services. This includes providing AI services to consumers and firms but also includes more directly providing financial services or facilitating AI-mediated interfaces for consumers to interact with retail financial services. The high cost of training AI models means that a limited number of AI model providers remain competitive at the frontier.⁴²

Secondly, AI can change how consumers purchase and interact with financial services.⁴³ New AI-mediated interfaces may be possible.

Thirdly, AI can change which retail financial services firms are the most competitive.

We explore each of these within this chapter. This analysis builds on existing competition authority work on digital markets, data-driven competition and AI foundation models, including concerns around control of key digital inputs, defaults, ranking, interoperability, switching costs and access to data or compute.⁴⁴

Competition across the AI supply chain will shape retail financial markets

One dynamic is barriers to switching AI models.⁴⁵ Technical or legal barriers to switching may increase the dependence on dominant providers. AI systems may learn from users and store contextual information as memory. This could improve performance but increase switching costs, creating lock-in for firms and consumers.⁴⁶

The AI model layer is not the only area where concentration is relevant. The diagram on the following page sets out some of the key parts of the AI supply chain in the context of retail financial services. AI chip providers and data centre providers are another source of concentration. Even if competition in the model layer is high, concentration in the provision of chips will affect retail financial market outcomes.⁴⁷

42 This was flagged by the CMA in 2023 in AI Foundation Models: Initial Report: https://assets.publishing.service.gov.uk/media/65081d3aa41cc300145612c0/Full_report_.pdf

43 For a broader discussion on the effects of AI on downstream competition, see OECD's "Artificial Intelligence and Competitive Dynamics in Downstream Markets" https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_c6e81d0e/ccf0624a-en.pdf

44 See, for example, Competition and Markets Authority, AI Foundation Models: Initial report, 18 September 2023, <https://www.gov.uk/government/publications/ai-foundation-models-initial-report>; European Commission, Competition in generative AI and virtual worlds: policy brief, 19 September 2024, <https://digital-strategy.ec.europa.eu/en/news/commission-publishes-policy-brief-competition-generative-ai-and-virtual-worlds>; OECD, Artificial Intelligence and Competitive Dynamics in Downstream Markets, 2025, https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_c6e81d0e/ccf0624a-en.pdf; Financial Conduct Authority, FS16/5: Call for Inputs on Big Data in retail general insurance, 21 September 2016, <https://www.fca.org.uk/publications/feedback-statements/fs16-5-call-inputs-big-data-retail-general-insurance>.

45 This was flagged by the CMA in 2023 in AI Foundation Models: Initial Report: https://assets.publishing.service.gov.uk/media/65081d3aa41cc300145612c0/Full_report_.pdf

46 Also discussed by DG Competition https://competition-policy.ec.europa.eu/document/download/c86d461f-062e-4dde-a662-15228d6ca385_en

47 For more discussion see OECD's "Competition in Artificial Intelligence Infrastructure" https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/competition-in-artificial-intelligence-infrastructure_69319aee/623d1874-en.pdf

Figure 10.
AI supply chain into UK retail financial services

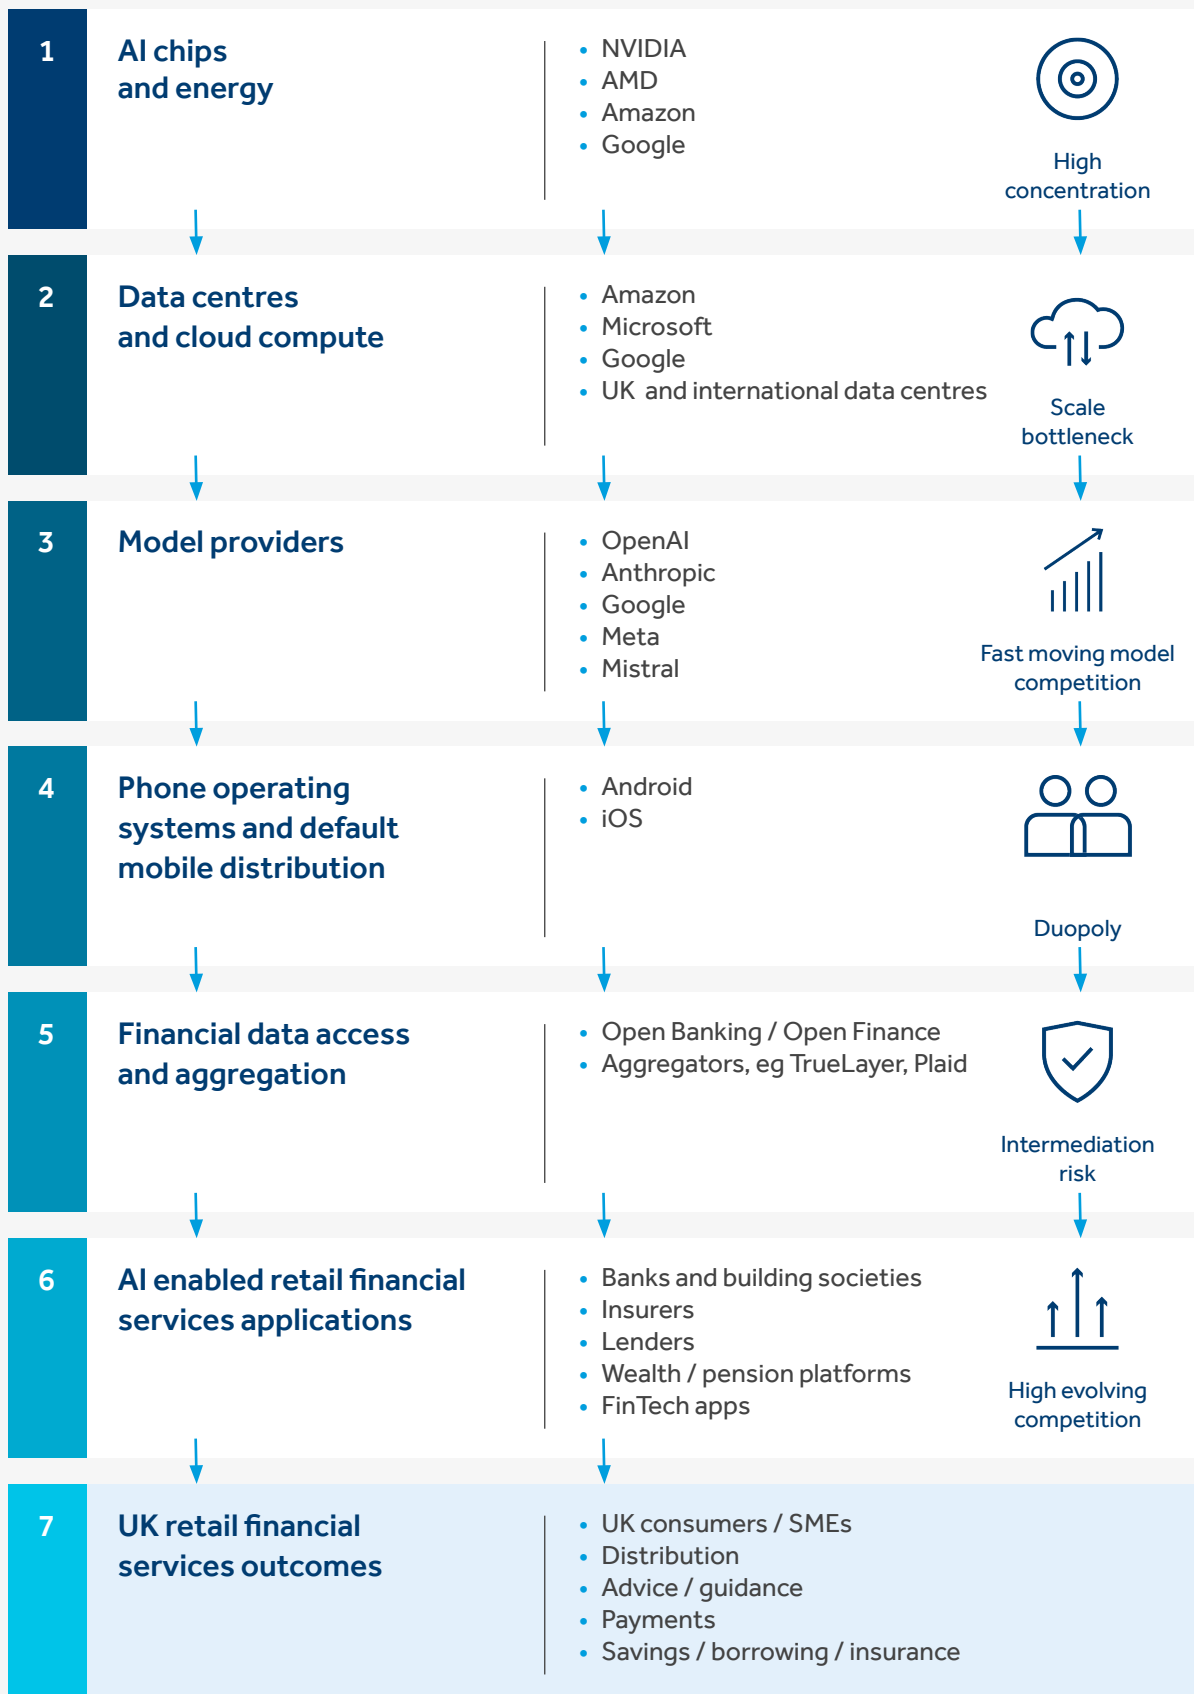


Figure 11.

AI-mediated consumer interfaces are likely to reshape competition and have the potential to impact market structure

At least 1 in 5 consumers used AI when shopping for current accounts or motor finance. Around a third used AI for pensions and investments, and nearly half among the minority seeking debt advice.

Chart showing % of consumers who use an AI-enabled tool to assist with financial products.



		% who shopped / used each in last 12 months	% using AI for each
1	Debt advice	2	45
2	Investing in shares, funds, other	10	34
3	Pension or retirement savings	5	29
4	Motor finance	4	25
5	Mortgages or other property finance	6	18
6	Other credit or loans	3	11
7	Pet insurance	9	15
8	Current accounts	19	21
9	Credit cards	19	18
10	Savings accounts	32	15
11	Home insurance	31	8
12	Vehicle insurance	37	8

F4. Thinking about the past 12 months, have you shopped around for, taken out, or switched any of the following financial products?

F5. Did you use AI to assist with or carry out tasks in any of these areas?

Base: all respondents (5026)

We discuss the regulatory implications of the level of concentration across the AI supply chain and its connection to regulatory regimes such as Operational Resilience and Critical Third Parties in the regulatory implications chapter and priority recommendation 2.

Turning to the second issue of consumer access, as consumers use AI services to inform, compare options and act, the interface that interprets their intent may become the point at which competition is shaped. That interface may decide which providers are visible, which products are compared, how trade-offs are explained, what defaults are used and when action is taken.

This creates a more complicated form of competition. Today, firms often compete for consumer attention through brand, app experience, branch presence, intermediary relationships, advertising, price, service quality and trust. In an AI-mediated journey, the consumer may not actively browse brands or compare providers directly. Instead, an AI interface may interpret the consumer's objective and narrow the options before the consumer sees them. That means a new layer of commercial value and competition in the supply of the AI-mediated consumer interface emerges.

Firms then have three broad options: they themselves may compete by supplying an agent or an interface, they might control or integrate with the interface, or they compete to be included, interpreted accurately, ranked favourably and selected by systems acting on consumers' behalf. Once they win or retain the customer or the customer's agent, they provide their traditional financial service: a loan, a savings account, a mortgage and so on.

The consumer may experience this as neutral assistance, but the interface may be shaped by commercial arrangements, technical integrations, subscription tiers, paid placement or other incentives. This will be explored in the different types of distribution models that could exist.

Of course, the consumer interface will not be AI-mediated for everyone by 2030. The consumer journeys chapter shows that demand for more autonomous AI is emerging but still limited. This means AI-mediated competition will coexist with more traditional forms of competition. Some consumers may actively prefer non-AI human-led routes, particularly for complex, high-value or emotionally significant decisions.






Stakeholders, especially intermediary firms such as wealth managers and brokers, have also pointed to continuing demand for human-led advice or brokerage services. There may therefore be competitive space for firms that deliberately offer trusted human-led or non-AI services, in the same way some consumers continue to value high-street branches or direct human support.

Different interface models create different competition dynamics

There is unlikely to be a single AI interface for financial services. Consumers could in principle access financial services through several routes, each with different implications for competition.

Possible interfaces include (see Figure 12. on next page):

Figure 12.
Five models for AI-mediated retail financial services

					
Phone OS controlled	AI app controlled	Bank app controlled	Third-party aggregator controlled	Customer chosen agent	
Customer ↓ Phone OS AI assistant ↓ Financial apps	Customer ↓ AI app ↓ APIs / integrators ↓ Financial services	Customer ↓ Bank AI agent ↓ First- and third-party financial services	Customer ↓ Aggregator platform ↓ Financial services	Customer ↓ Financial services	AI control
Phone OS	AI app	Bank app	Aggregator app	Customer provided agent	Main interface
OS mediated	AI firm mediated	Bank led	Multi provider	Agent led	Provider choice
Linked providers	Linked providers	Bank products	Linked providers	Linked providers	Funds source
Platform gate keeping	Platform gate keeping	Incumbency / walled garden	Data access / interoperability	Data access/ interoperability	Illustrative competition concern
Mainly financial provider; OS may share platform liability	Financial provider; AI applicable for its own intermediation	Bank	Aggregator for routing; provider for product execution	Consumer for routing; provider for product execution	Possible liability holder

The Phone Operating system model could scale quickly because it sits close to everyday digital behaviour and may become the default route through which consumers use AI.

Historical operating system level financial integrations, such as Apple Pay and Google Pay, demonstrate how rapidly embedded services can gain consumer adoption at scale. If consumers ask an embedded AI agent to manage payments, compare financial products, or identify an insurance product, the financial firms' own brand and agent interfaces may become less central. Competition could then depend increasingly on firms' access to the OS ecosystem. The competition risk is therefore platform gatekeeping: the operating system provider may control access, defaults, ranking and commercial terms.

General-purpose AI apps have an advantage in that many consumers already use them to ask financial questions: 28% of existing UK AI users and 16% of all UK adults, which equates to approximately 9 million UK adults. Some even initiate product selection through integration with financial firms. This could create a new cross-market gateway into financial services, especially where consumers already trust or habitually use the AI app. The opportunity is wider access to financial information and support; the risk is that financial influence may sit outside regulated financial services, particularly where outputs resemble guidance or advice.

This model could create competitive tension between regulated firms and out-of-perimeter AI platforms. A general-purpose AI service may shape consumer understanding, narrow options or recommend courses of action without being part of the traditional financial services relationship. That could create an uneven playing field if regulated firms face obligations around communications, advice, promotions and consumer outcomes while AI platforms exert similar influence through a different route.

Banking apps may be enhanced with AI functionality. This model allows incumbents to use existing trust, data, brand and customer relationships to defend their market position. A bank could improve customer support, personalise services, prompt action, integrate products and make its own app more useful, and through Open Finance (discussed in Annex VII) eventually offer a whole market view.

Enhancing these apps with AI could improve service and strengthen trust, but it may also limit switching if the AI interface is designed around retention rather than genuine whole-of-market comparison. A provider app has different incentives from an independent agent. A bank may be able to build an AI application that helps consumers manage money, but it may have weaker incentives to encourage movement to competitors, unless required by regulation, competition or consumer expectations. In addition to banks, this approach could also be undertaken by an insurer or investment platform.

Third-party aggregators may integrate AI which can search, compare and execute across multiple providers. This could intensify comparison and preserve a role for intermediaries if they offer trusted market coverage, explanation and execution. It may also help consumers understand more complex trade-offs and reduce the effort required to switch.

The risk is that comparison may appear broader than it is. Commercial arrangements, paid placement, incomplete integration and ranking criteria could shape what the consumer sees. AI may make this less visible because the consumer may receive a conversational answer rather than a transparent list of included and excluded providers.

The customer choosing an independent AI agent is another approach, which may support or act for them across providers and accounts within permissions the customer sets. This could be the most consumer-led model if the agent can search broadly and act in the consumer's interests across markets. It could reduce inertia, prompt switching and make competition more active by enabling consumers to act without repeatedly comparing and authorising each step themselves.

But this model could depend most heavily on the enabling infrastructure: data access, interoperability, identity, mandates, liability and audit. Without those foundations, independent AI agents may remain partial, fragmented or unsafe.

Across all of these models, by 2030, we anticipate that the capability, commercial and consumer impact of these AI-mediated consumer interfaces in financial services could be accelerated by certain infrastructure enablers. These include open finance, programmable digital money through the UK's National Payments Vision including stablecoins, tokenisation, and potentially central bank digital currencies. They could broaden the data available to agents, make products and transactions easier to manage, automate and execute in many areas. Mortgages are an early example for tokenisation-type solutions.

From a competition perspective, they will increase the incentives for firms to control the consumer interface layer in financial services. The competition chapter links directly to the report's later recommendation on agentic finance and Annex VII analyses data, identity, delegation, payments, liability and supervision.




Interfaces can connect to financial services through different integration models

AI agents could connect to financial services in three broad ways (see diagram on next page). These are: direct bilateral deals between AI platforms and retail financial services firms; private gateways covering parts of the market, or open standards, such as Open Finance.

These are different models that may develop in parallel or even compete with each other. The importance for competition (and for the regulator) is that the balance between them might affect market outcomes, such as contestability and consumer choice, depending on whether firms rely on direct deals with selected firms, connect through private intermediaries, or operate through open and standardised access arrangements.

Direct deals between AI platforms and retail financial services firms require each financial firm that integrates with the AI model provider to opt in, but integration with a single financial intermediary, such as a price comparison website or investment platform, can make multiple - but not all - financial products available to the customer.

Figure 13.
Three models for AI integration into financial services

 <p>Direct deals</p> <p>Access depends on one to one deals or custom integrations between specific firms and particular AI platforms.</p> <p>Examples</p> <ul style="list-style-type: none"> Firm specific AI assistant integrations Specific insurer or price comparison app inside an AI assistant Custom bank or insurer connectors 	 <p>Private gateways</p> <p>A private intermediary connects to many firms, giving broad reach. Access, ranking, fees and participation are governed privately.</p> <p>Examples</p> <ul style="list-style-type: none"> Banking data aggregator Price comparison websites Aggregator and marketplace models 	 <p>Open standards</p> <p>Access is based on common rules and interfaces. In principle any qualifying firm or agent can connect on similar terms.</p> <p>Examples</p> <ul style="list-style-type: none"> Open Banking Open Finance (future and emerging) Standardised permissioned data access 	
<ul style="list-style-type: none"> Fast for individual firms Can be very polished Clear commercial incentives 	<ul style="list-style-type: none"> Broad coverage Fast scaling Convenient discovery 	<ul style="list-style-type: none"> Broad market access Portability and switching Less dependence on one platform 	Benefits
<ul style="list-style-type: none"> Patchy coverage Fragmentation Incumbent advantage and deal making power 	<ul style="list-style-type: none"> Gatekeeper power Ranking and default bias Private control of access 	<ul style="list-style-type: none"> Standards may evolve slowly Access and liability rules matter Consent UX still matters 	Risks
Narrow	Broad	Broad	Breadth
Private / bilateral	Private	Open / standardised	Governance

There has been some early uptake of financial firms integrating with specific AI interfaces already, including by banks, insurers, and price comparison websites.^{48 49 50} Additionally, AI platforms have started to enable users to connect financial accounts using data aggregation providers.⁵¹

For **open standards**, Open Banking and Open Finance are relevant because the standards they set could inform how financial services firms open their products to AI agents, in a way which is comparable. One advantage of standards over specific financial firm integration is broader cross-firm visibility, which would be more powerful for comparison, advice and whole-of-market functionality and likely support more growth and innovation in financial services markets.

The challenge will be to move Open Finance forward at pace given the strong commercial incentives for all firms, regulated and non-regulated, to capture value from AI interfaces. Another route is for private financial services intermediaries with broad (but not complete) financial market coverage, which we term **private gateways**.

This distinction is important. The consumer may see only one AI interface, but the breadth, neutrality and contestability of the comparison will depend on the integration model underneath it.

Firms will compete within AI-mediated channels

Financial services products can be distributed through new interfaces set out above. Firms may compete across more than one channel, and within each channel compete for visibility, trust and conversion.

Control may sit at several different layers of the consumer journey: the provider of the underlying financial service, the provider of search or comparison, the provider of the agentic model, or the holder of customer data.

The nature and effect of these new distribution channels relate to the spectrum of autonomy outlined earlier. Today, customers operating financial services can collaborate with AI, but the effect on competition is limited by patchy access to customer data and the inability of most agents to take actions. By 2030, competition between interface models could intensify if these access and execution capability challenges are no longer a factor, moving financial services down the autonomy spectrum towards humans approving AI prompts, or observing AI actions. This is covered in the agentic finance infrastructure annex.

Faced with AI-mediated channels, financial firms will respond so their products and services are more likely to be selected, which could include embedding AI internally so that they remain good value or by making their products and services easier for agents to access and compare.

48 <https://www.natwestgroup.com/news-and-insights/news-room/press-releases/ai-and-data/2026/apr/natwest-becomes-first-uk-bank-to-launch-trusted-home-buying-adv.html>

49 <https://www.aviva.com/newsroom/news-releases/2026/03/aviva-launches-insurance-app-on-chatgpt/>

50 <https://www.moneysupermarket.com/news/new-moneysupermarket-chatgpt-app/>

51 OpenAI announced in May 2026 the ability for users to connect financial accounts, initially for US users and through Plaid, <https://openai.com/index/personal-finance-chatgpt/>

Others may focus on customers who do not use AI interfaces or improve the AI-enabled user experience in their own apps, using trust, existing relationships, and integrated servicing to keep customers within direct channels. Finally, some may aim to build their own AI-enabled interfaces.

AI may shift the nature of intermediation towards integrating with agents, platforms, and data-access

Retail financial services markets are highly intermediated. In general insurance, only around 30% and 20% of policies are sold directly for home and motor respectively, while around 90% of first-time buyers purchased mortgages via a broker.^{52 53}

AI challenges intermediation. If customers can use agentic AI to search and compare more of the market with lower effort, the information gap between firms, intermediaries and customers may narrow.

As customers increasingly use AI assistants or agents to search and compare and select products, firms may need to present product information and regulated communications in ways that are legible to those systems as well as to humans, within the context of regulations on financial promotions.

Where intermediation reflects differences in distribution or onboarding costs, AI could make it more attractive for product providers to serve customers through their own channels rather than relying on intermediaries.

However, one key driver of intermediation is trust, and in particular having a specific human to speak to throughout a complex and high value process such as getting a mortgage. Our research shows that customers trust AI more readily for lower value financial services, or where the consequences of mistakes are lower.

Furthermore, even in a world of widespread use of AI in financial services, many customers will still prefer to deal with intermediaries, either because this is how the customers have done business in the past, or because customers prefer the experience and therefore trust it or because the way AI consumer interfaces develop does not provide consumers with whole of market comparison and they may prefer intermediaries who can provide that (albeit, they will likely have AI as a core operational capability supporting human-led service). In addition, regulations in some sectors require customers to seek advice.

AI delegation may not remove intermediation so much as relocate it: from brokers, comparison tools or product providers towards AI agents, platforms and data access layers that shape what consumers see and can act on.

52 <https://www.fca.org.uk/publication/corporate/ep25-2.pdf>

53 FCA 2023 PSD data <https://www.fca.org.uk/publication/consultation/cp25-11.pdf>

AI interfaces may become a new competitive gatekeeper and may become a source of market power

AI channels will compete with existing traditional routes to market such as advertising, comparison websites, search, brokers, and direct. The commercialisation of these new routes to market is emerging and might comprise integrations with model providers (we see these with OpenAI's ChatGPT already), paid placements, better integrations or APIs, exclusive partnerships or trusted-agent certification processes or protocols. In this sense, competition might move upstream.

Competition may still occur through price, as with comparison websites, but also through defaults and convenience, as seen with Apple Pay and Google Pay

Experience with digital wallets suggests that, when a viable interface business model emerges, firms may compete strongly to capture it. The regulatory challenge is to preserve openness and contestability of retail financial markets without preventing viable AI-mediated channels from developing.

We could see impacts on the structure of competition if AI-directed flows may favour mainstream, visible or easily integrated providers, increasing concentration in retail financial services if products from smaller or specialist firms are less likely to be presented to the customer. As a result, even where the same products remain available, the effective choice set may narrow.

The challenge will be whether effective comparison and choice is preserved. There is a risk that AI interfaces could steer users through a narrower or less transparent subset of options depending on commercial arrangements, subscription tiers or paid access models. In effect, AI-mediated comparison could appear to search the market while in practice the results are narrowed by technical integrations, commercial arrangements or paid prioritisation.

Firms may also compete by making it easier and safer for agents to act for consumers. This could include verifying customer identity, giving agents secure access to customer data and managing agent permissions when connecting with financial services (see the analysis in the agentic finance infrastructure annex and our priority recommendation 5).

AI interfaces may drive hyper-switching

Many stakeholders mentioned the potential for AI significantly increasing the switching of financial products and services by consumers. This includes reducing the costs of switching and the informational barriers to switching. They also mentioned AI agents enabling continuous personal financial management and execution leading to automatic and rapid switching and portfolio management in the approver and observer stages on the autonomy spectrum.

Stakeholders envisaged this being a pro-competitive force across insurance, investment platforms and retail banks, though one that could erode margins on traditional products for firms and potentially spur changes to product design and pricing.

Some stakeholders raised the risk in retail banking that in a world of empowered AI agents, able to shift customer savings between banks and building societies instantly and effortlessly, the banking market could be more fragile, leading potentially to financial stability implications.

If deposits are easier to move, it could affect banks' funding models and costs of capital. In turn, the cost of lending for consumers, small businesses and corporates could increase to compensate for lower net interest margins, and the structure of competition might favour those with a wider variety of capital, funding and income sources than simpler business models. Over time, these developments could lead to consolidation or the growth of more diversified business models.

Firms across retail financial services might need to seek to differentiate or compete to retain customers and their agents. They may respond by strengthening retention offerings, more bundling and new loyalty schemes or other product features.

There are some areas of the UK financial services system where models are based on member interests, such as, building societies, mutuals and credit unions. These models can benefit from AI, if they adopt this, by finding more efficiency in their services. But they will need to change, modernise and adopt AI and digital technology to support their members' interests in the future. There are risks to many business models from both upstream dependency and downstream competition at the consumer interface. One specific risk that arises might be if agents are able to move money rapidly such as deposits. Those businesses more reliant on deposit taking activities to fund other parts of their book might need to transition to continue to meet their customer's needs. We have not assessed the extent of potential consolidation or whether there is a need to underpin wider activity in these sectors, but we do consider that financial regulators should pay specific attention to the impact of advanced AI on mutual and credit union business models to ensure that this critical part of the financial services sector can transition sustainably over time.

AI interfaces may shift competition towards the features the interfaces can access most easily

Customers using AI interfaces to compare financial products may be better able to compare prices and some non-price features, potentially increasing competition on both price and product quality in affected retail financial services markets. Some stakeholders suggested that firms whose business models are built on consumer habit, information asymmetry, or inertia, rather than genuine value creation, face the most acute disruption risk.

Customers benefit from lower prices, however the FCA has previously explored whether increases in price competition cause firms to introduce lower quality products with lower costs of provision.⁵⁴

⁵⁴ For example, the FCA's "Evaluation Paper 25/2: Our general insurance pricing practices (GIPP) remedies" which explored the impact of pricing remedies on product quality. <https://www.fca.org.uk/publications/corporate-documents/evaluation-paper-25-2-general-insurance-pricing-practices-remedies>

The effect will in part depend on how well the AI can assess these non-pricing factors, and how the AI interfaces present choices to customers. The dynamic could even go the other way, with less effective AI interfaces placing more weighting on non-pricing factors.

AI may blur competitive boundaries around the perimeter

Industry engagement highlighted concerns about a level playing field. Retail financial services are often delivered as bundles of regulated and adjacent activities. General-purpose AI providers may be able to unbundle some parts, such as information, support or comparison, from outside the perimeter, before later monetising through referrals, commissions or deeper partnerships.

Competition between activities inside and outside the perimeter may therefore become more complex, creating risks of arbitrage, uneven accountability and market distortion.

Much of this dynamic depends on AI labs and other adjacent providers remaining outside the perimeter. If some of these firms later seek authorisation, partner more deeply with regulated firms, or regulated activities move inside the perimeter, competition could shift again, because firms that built scale, customer reach or data advantages outside the perimeter could then enter regulated markets from a stronger position. We discuss the need for the FCA to examine the impact of AI on the regulatory perimeter in priority recommendation 1.

AI-enabled firm capability will become a core driver of competition, reshaping how incumbents, challengers, smaller firms and new entrants compete

The firms chapter considered how firms are embedding AI into their operating models, this section considers the same shift from a competition perspective: which types of firms are best placed to turn AI adoption into a competitive advantage.

AI may help existing smaller firms compete more effectively by reducing operating costs and giving access to capabilities that previously required larger teams. However, benefits are likely to be uneven: smaller firms may have less proprietary data, fewer technical staff, weaker vendor bargaining power and higher relative governance, testing and assurance costs.⁵⁵

We heard from some smaller firms that they can benefit from the early adoption and expenditure of larger firms, as they then see more competition flow once vendors have worked through efficient and effective solutions for larger firms. They also commented that the lag between large firm adoption of new capabilities and movement through to smaller firms has shortened.

Across digital native and smaller firms, a key competition enabler is likely to be the continued surge in AI capability that we describe in our system driver chapter.

⁵⁵ FCA, Understanding the advice market: financial advice firms survey 2025. The survey found that 95% of large advice firms were using, or considering using, AI in the advice process, compared with 48% of small firms. <https://www.fca.org.uk/data/understanding-financial-advice-market>

And, as we set out in the firms chapter, the ability of firms to embed these capabilities into their operations successfully will likely distinguish winners and losers in the future.

That said, our research shows that in the run-up to 2030, incumbents may still benefit from proprietary data, brand and customer trust, capital, governance and vendor relationships. Larger firms may also be better placed to negotiate access, pricing and support from AI and market-specific vendors. A key indicator is the recent release of Mythos, where only the very largest incumbents were able to both gain access and afford to test their systems against these new capabilities.

Deploying AI effectively will not mean just adopting generic tools. Firms need tailored implementation, testing, monitoring and integration with existing workflows. This favours those with proprietary data, technical capability and capacity to absorb upfront costs.

High fixed costs, skills shortages and infrastructure constraints may also favour larger firms and limit access for smaller firms or consumers. And sovereignty and security issues might also delay frontier AI capability reaching UK firms and this could also impact competition among incumbents.

Finally, AI may also lower some barriers to entry. AI-native entrants may avoid legacy technology constraints and design workflows, data architecture and controls around AI from the start.

Entrants may also be able to do more with off-the-shelf tools than more established firms, reducing the cost of testing propositions and scaling some functions. But entry will still depend on authorisation, trust, distribution and access to market-specific data.

We have engaged with firms who have entered financial services markets in non-UK jurisdictions as digital AI-native firms and have gained share. They demonstrate that the speed and pace of their ability to complete operational, front office and risk and compliance activity powers their ability to gain share from traditional finance competitors. These are signals that AI-native firms might be able to compete, and over time, gain share.

AI may therefore intensify concentration in some markets

Competitive pressures may also lead to exit by firms that cannot adapt, particularly where AI raises the minimum efficient scale of operating or distribution.

Customers can still benefit even if existing firms gain share, for example if effective AI deployment lowers prices, improves service quality or expands access. In some markets, firms that can deploy AI effectively at scale may lower costs, improve quality or roll out better products more quickly, even if that also puts pressure on smaller competitors.

AI may intensify consolidation and vertical integration in some markets, particularly where it strengthens scale, bargaining power, data advantages or distribution control. Some AI-native firms may be subsequently acquired, smaller firms may consolidate to compete, and larger firms may integrate across products, data, models and distribution.

Of course, scale, trust and data advantages do not automatically imply mass consolidation. Financial services firms already benefit from economies of scale, shared data infrastructures and strong brands in many markets. AI may add to these existing dynamics without necessarily leading to mergers or highly concentrated outcomes everywhere.

Implications for the rest of the report

AI will reshape competition in retail financial services through three connected layers. First, competition in the upstream AI supply chain will determine the terms on which financial services firms can access the models, compute, cloud infrastructure, data and specialist vendors needed to deploy AI. If these markets remain contestable, firms may benefit from better capability, lower prices and faster innovation. If they concentrate, firms may face higher costs, restricted access, weaker bargaining power and less innovation in the AI tools on which they depend.

Second, as consumers use AI to search, compare and eventually act, market power may shift towards the AI-mediated interfaces that interpret consumer intent and determine which products are visible, comparable and executable.

Third, firms will compete on how effectively they embed AI into their own operations, lowering costs, improving service, strengthening governance and building trust. The impact will vary across incumbents, smaller firms, digital challengers and AI-native entrants.

These layers reinforce each other. Competition in the AI supply chain shapes the capability available to firms and interfaces; AI-mediated interfaces reshape how firms compete for access to consumers; and firm-level AI capability determines which firms can outcompete rivals.

The competitive landscape will therefore be shaped not only by competition between financial services firms, but also by competition within the technology markets that increasingly underpin them.

System shift 4

Amplified financial crime and cyber risk

System shift 4: Amplified financial crime and cyber risk

Key Findings

- AI will amplify fraud and cyber risks by 2030: they will become faster, cheaper, more scalable and more persuasive.
- The main pressure is speed and scale. Faster exploitation of existing weaknesses, whether customer deception or weak cyber controls.
- Risks will cut across the system be it firms, platforms, telecoms, payment rails, identity systems and AI/technology providers and jurisdictions.
- AI is dual use; it strengthens detection and resilience but needs governance and oversight.
- Capability must evolve at system level, across intelligence-led prevention, coordination, disruption, cyber fundamentals and clear escalation routes.

Engagement paper responses

- Fraud and financial crime threats are growing rapidly and becoming more complex.
- Respondents said AI is making fraud and scams quicker to spread and more persuasive, especially impersonation, synthetic identity abuse and automated social-engineering attacks.
- Responses were concerned about the impact and rise of AI-enabled harms, and there was uncertainty about the FCA's enforcement approach to new AI risks. Evidence from this review showed a clear shift in the fraud threat by 2030.
- Responses indicated uncertainty about the regulatory response which may contribute to inconsistent adoption of AI-enabled fraud and cyber controls, as firms balance the need for faster detection and disruption against expectations on governance, accountability, human oversight and customer harm.
- There is potential for a capacity gap which criminals can exploit. But responses also recognised the opportunity and benefit of coordination to share real-time data, emerging typologies, and coordinate action. This fragmentation is compounded where fraud operates across borders, with differing regulatory frameworks, data-sharing constraints and response speeds.
- According to external evidence, social platforms are compounding this fragmentation, with breached personal data and seemingly benign interactions that build trust before subsequent exploitation.
- Internal material and wider FCA work point in the same direction: harm will spread across onboarding journeys, payment rails, telecoms and digital platforms, rather than sit within a single jurisdiction or firms' visibility or control.

Capability pressures in the AI-enabled fraud landscape

AI is more likely to accelerate existing fraud typologies than create a wholly new category of crime. Better tools, industrialised criminal infrastructure and cross-border ecosystems allow tactics to be deployed, tested and refined at scale.

The 2025 UK National Risk Assessment identifies fraud as a major money-laundering threat and AI as a cross-cutting risk shaping future threat patterns.⁵⁶ Analysis by the AI Security Institute highlights the growing role of AI in cyber-enabled fraud and exploitation of system vulnerabilities.⁵⁷ The Association of certified Anti-Money Laundering Specialists (ACAMS) has described AI-powered identity fraud and scams as among the most significant global threats to financial crime prevention.⁵⁸ By 2030, this could mean more attempted fraud and greater pressure on systems to detect, escalate and disrupt harm. Stakeholder evidence and wider reporting echo this.

AI is lowering the barriers to fraud. Techniques such as cloned voices, synthetic identities and highly realistic fake content are making scams easier to produce, cheaper to scale and harder to spot. This enables a much wider pool of criminals to carry out convincing fraud, especially when combined with compromised personal data across onboarding, payments, telecoms and digital platforms. As the same time, by 2030 consumers may rely on AI agents to act for them. As people shift from consultant to approvers of AI agents - whether these are mainstream consumer platforms or use of independent AI agents - as outlined in the consumer journeys chapter, consumers may stop challenging decisions made on their behalf or over rely on AI applications. This combination will amplify financial crime and cyber risks.

Put simply, AI lowers the cost of reconnaissance, personalisation and iteration for criminals, while increasing the burden on firms to distinguish genuine customers from convincing deception. Firms may optimise for known risks and work on AI-enabled sanctions evasion points to this challenge: many institutions may remain around current typologies while adversaries are adapting faster, including using international transaction pathways and distributed infrastructure and subsequently miss new risks; regulators may receive more reports but less timely intelligence.⁵⁹

As these pressures intensify, lower-risk detection and triage may increasingly sit at AI agent levels of delegation. Humans will still need to shape tasks, review outputs and judge exceptions, but a proactive approach will need stronger human oversight, given the stakes for fairness, accountability and customer harm.

⁵⁶ In section 6 - cross cutting risks the UK's National Risk Assessment of Money Laundering and Terrorist Financing 2025 provides analysis of the emergence of AI as both a capability builder for the prevention of money laundering and terrorist financing and for exploitation by criminals. See [National Risk Assessment of Money Laundering and Terrorist Financing 2025](#).

⁵⁷ The AI Security Institute paper presents analysis of the trends in frontier AI. See AI Security Institute, [Frontier AI Trends Report](#), 2025.

⁵⁸ The ACAMS third annual report, the Global Anti-Financial Crime Threats Report 2026, forecasts the top 10 threats for 2026. This includes technological threat with AI having a reshaping influence. See [ACAMS Global AFC Threats Report 2026](#)

⁵⁹ The Royal United Services Institute (RUSI), [Beware the Robots: AI-Enabled Sanctions Evasion is Here](#)

System resilience and cross-sector capability

Figure 14. (next page) illustrates how AI-enabled harm may move across firms, platforms and infrastructure by 2030, increasing the need for system-level coordination.

By 2030, fraud risk is likely to become more fragmented and cross-sectoral and increasingly international in nature. Even where individual firms strengthen controls, harm may move through the links between banks, fintechs, telecoms providers, online platforms, identity systems and payment rails.

Criminals can shift across these channels faster than institutions and authorities can coordinate. Capability therefore needs to evolve at system level, across firms, sectors and public authorities, particularly where no single organisation has the visibility, data or operational reach to act alone.

This could include pooled specialist resources, shared analytics and more collective approaches to disruption, with criminals treated as the common adversary.

Respondents pointed to three recurring frictions: slow and inconsistent intelligence-sharing, ability to detect genuinely new typologies, and system responses that arrive too late to prevent harm.

The challenge is around combining and acting on signal data. The call for evidence on economic information sharing acknowledges the 'UK's economic crime information sharing system remains a complex patchwork'.⁶⁰ The Financial Data Transparency Act (FDTA) in the US has also taken steps to finalise data standards between regulators in a bid to move to standardised, machine-readable interoperable data⁶¹. Inconsistent data standards and quality, uneven identifiers, and uncertainty over lawful information-sharing routes have the potential to delay action. This challenge is amplified in cross-border contexts, where legal gateways and standards vary.

Sharing data, stronger identity attributes and risk indicators could improve detection and reduce the risk that false legitimacy builds over time. But these depend on trusted governance, interoperable standards and confidence that data is used lawfully and proportionately.

Recent developments - including online crime coordination, legal and data-sharing enablers, privacy-conscious data initiatives and Economic Crime Plan 2 - suggest the foundations for faster collective response are emerging, but it is still uncertain whether those foundations can be connected, scaled and operationalised quickly enough.

A credible system will need a minimum set of signals. This means defining the core warning signs that firms, regulators and partners should all collect and share in a consistent way, so emerging harm can be spotted, escalated and disrupted earlier.

⁶⁰ The UK Government call for evidence on economic crime information sharing in March 2026 is looking at, despite the progress made, the UK's economic crime information sharing system remains complex. See, [Economic crime information sharing](#), 2026.

⁶¹ Office of the Comptroller of the Currency, OCC Bulletin, [OCC Bulletin](#), 2026.

Figure 14.
Where the system will strain first: future pressure points

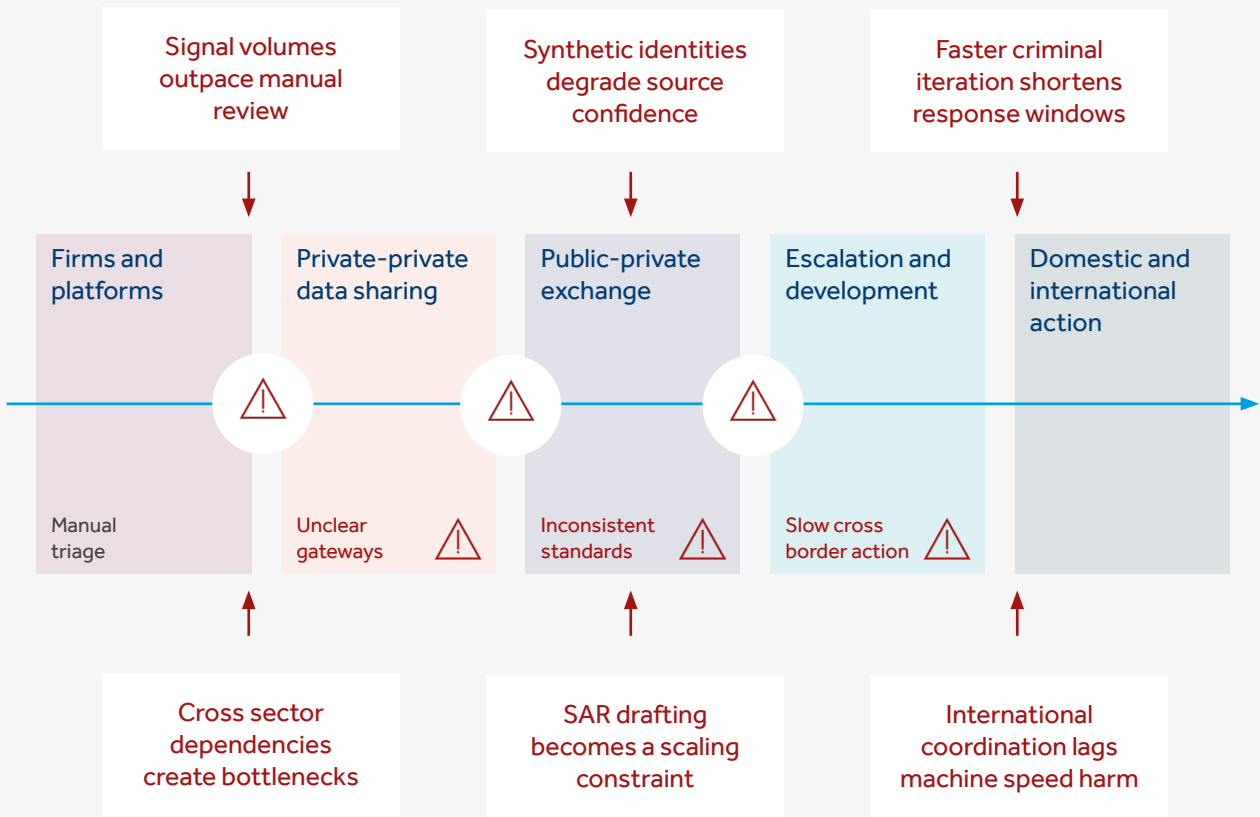
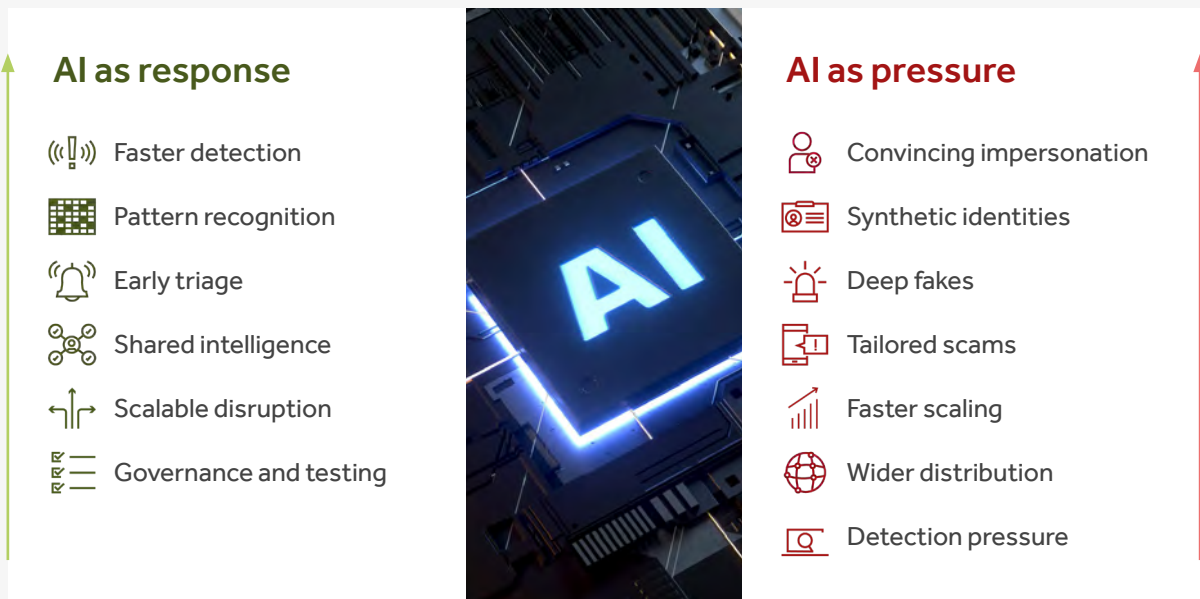


Figure 15.
How AI-enabled systems may act as both accelerants of harm and components of defence, depending on governance, delegation and oversight



To work in practice, this will require common data standards, clear information-sharing gateways, interoperability across FCA systems, firm reporting and partner data, and alerting that are directly linked to escalation and action.

Without those foundations, the system risks producing more analysis without accelerating disruption.

Evolving capability: AI as both pressure and response

By 2030, parts of the fraud response may need to operate at machine speed. As discussed in the firms chapter, AI models are already enhancing real time monitoring and pattern recognition in fraud defence.⁶² And AI-enabled systems could further help firms and regulators identify weak signals and connect patterns across data sources earlier including signals that span jurisdictions and sit across multiple regulatory or geographic boundaries.⁶³

As capability moves along the autonomy spectrum, the role of humans in detecting and managing fraud and financial crime is likely to shift. At lower levels of autonomy, AI assistants may support human operators by surfacing signals, identifying patterns and prioritising alerts, with humans retaining direct control over investigation and decision-making. As systems take on a more advanced role AI assistants may undertake initial screening, triage and the connection of signals across datasets, with humans focusing on validating outputs, reviewing exceptions and determining escalation.

At AI agent level systems may act with greater independence, human involvement is likely to concentrate on oversight, setting parameters, and making high risk or more consequential decisions, particularly where judgement, accountability or customer impact are significant.

The same capability can cut both ways: poorly governed AI may create false assurance or amplify risk, while well-governed AI could strengthen earlier detection, triage and disruption.

Figure 15. (previous page) shows how AI-enabled systems may act as both accelerants of harm and components of defence, depending on governance, delegation and oversight.

Using AI without creating new blind spots will be challenging. Automation can improve speed, but without strong testing, unclear accountability and effective oversight, it risks overwhelming teams with non-actionable alerts and could mean they miss novel behaviour. By 2030, agent to agent interactions is likely to add a new dimension to fraud defences. As AI systems interact directly fraud may occur within automated processes rather than at human interface. Defences will need to adapt to detect and intervene in these machine-driven interactions.

Firms should be able to show that AI-enabled controls improve detection, triage or disruption, remain effective as threats evolve, and retain meaningful human control where decisions affect fairness, accountability or customer harm.

62 European Central Bank, [AI's impact on banking: use cases for credit scoring and fraud detection](#), 2025.

63 Yang, H.; Shukur, Z.; Sahran, S. [A Review of Artificial Intelligence for Financial Fraud Detection](#), 2026.

Building effective fraud defence

By 2030, effectiveness will hinge on whether firms' use of AI in fraud controls demonstrably strengthens their defence against faster, cheaper and more adaptive fraud and novel risks.

This requires deploying controls at pace with sufficient assurance, governance and coordination, supported by frameworks such as the Senior Managers Regime and model risk management.

Cross-firm, cross-sector and cross border intelligence sharing will detect patterns earlier and enable a faster response to a system-wide threat. These shifts involve trade-offs between certainty and flexibility, and between speed and completeness of information. But as generative AI reduces the cost, speed and realism of fraud, firms' defensive capability must equally adapt.

Operational implications for the FCA

If effective capability by 2030 depends on faster coordination, usable signals and earlier disruption, the FCA will need to join up authorisations, supervision, intelligence and enforcement more closely.

Early warning signs may appear quickly and from many places at once, including firm reporting, supervisory work, whistleblowing, consumer complaints, transaction data, partner agencies, law enforcement, online platforms and other external intelligence including international partners. The FCA will need to bring these inputs together quickly, decide who is best placed to act, and coordinate a response before harm spreads.

In practice, harm may arise in one part of the system while the ability to assess or disrupt it sits elsewhere. Early signals may not reach those best placed to act. The FCA will need clearer pathways to connect detection, assessment and escalation. Without this, signals may remain isolated rather than acted on at pace, including where coordination is required beyond domestic boundaries. An agentic supervision model would support this shift (priority recommendation 6). AI-enabled systems could support lower-risk triage, but interventions should retain human oversight, given the stakes for fairness, accountability, explainability and customer harm.

The FCA will also need better shared signals, interoperable data and clearer escalation routes across firms, sectors and public authorities, strengthen system-wide coordination and oversight (priority recommendation 2).

More data alone will not be enough: *signals must be intelligible, linked to action and capable of supporting scalable disruption.*

As the Online Crime Centre, Economic Crime Data Strategy and information-sharing reforms develop, the FCA will need to work through shared standards, public-private coordination, and faster disruption mechanisms in addition to its supervisory and enforcement tools.

Cybersecurity

Increasing AI capability has deep implications for cybersecurity. In testing Anthropic's latest frontier model, the AI Security Institute (AISI) has shown that it 'can exploit systems with weak security posture, and it is likely that more models with these capabilities will be developed.' This assessment is supported by the Five Eyes National Cyber Security Centre statement recommending organisations to act on rapidly transforming cyber risks.⁶⁴

However, it explicitly recognises that 'AI cyber capabilities are dual use; while they pose security challenges, they can also help deliver game-changing improvements in defence'.⁶⁵

Frontier AI is a step change in AI capabilities, with significant implications for cyber security and operational resilience. Firms that have under-invested in fundamental cyber security are likely to become progressively more exposed.

In 2026, Anthropic's Mythos Preview model acts as a signal. Anthropic describes it as a 'watershed' moment for cybersecurity as it can identify and exploit previously undiscovered vulnerabilities in real-world software, among other claims about its ability.

Some of these have been independently corroborated by institutions such as AISI.⁶⁶ However, it is worth bearing in mind AISI's caveat that 'we cannot say for sure whether Mythos Preview would be able to attack well-defended systems'.

Bain and Company has said that the era of AI-powered attacks at scale has arrived, and that AI does not create new vulnerabilities, it exposes existing ones.⁶⁷ Organisations' call for strengthening cybersecurity fundamentals is an immediate priority, since attackers only need to find one flaw.

Cybersecurity in 2030

Mythos Preview is not the only frontier model where cybersecurity capabilities are advancing.⁶⁸ Therefore, it is reasonable to assume that this restricted capability will become widely available by 2030. If the ability to find and exploit software flaws becomes available at scale it creates a significant pressure on organisations to triage and fix flaws at sufficient speed.

This could be amplified by the shift along the autonomy spectrum. Autonomous AI agents may be able to plan and execute attacks end-to-end with little or no human direction (as they are an observer). Further, these attacks are likely to be able to be carried out at machine speed, so defence will have to be automated to match.

64 G7 Cyber Expert Group, [Statement on Artificial Intelligence and Cybersecurity](#), 2025.

65 AI Security Institute (AISI), [Our Evaluation of Claude Mythos Preview's cyber capabilities](#), 2026.

66 AI Security Institute (AISI), [Our Evaluation of Claude Mythos Preview's cyber capabilities](#), 2026.

67 Bain and Company, [Claude Mythos and the AI Cybersecurity Wake-Up Call](#), 2026.

68 Bain and Company, [Claude Mythos and the AI Cybersecurity Wake-Up Call](#), 2026.

The scale and impact of such cyber-attacks may be greater because by 2030 much critical financial services infrastructure will run on AI. This increases the attack surface for the threat actors described in this section.

There is also a potential risk from upstream concentration. As the CEG states, 'a significant cyber incident at a widely used AI provider could have the potential to affect many financial institutions'.⁶⁹

The cybersecurity implications also change as the human role changes. At lower levels, AI helps human attackers or defenders find information, write code, analyse vulnerabilities or prioritise action. At higher levels, AI-enabled systems may coordinate attack chains or defensive responses with less real-time human direction.

This increases the premium on clear controls: firms will need to know which cyber decisions can be automated, which require human approval, and which require active human judgement because of their operational or customer impact.

There is the potential for this pressure to intensify through quantum computing (discussed previously in this report) which is expected to threaten much of the encryption that currently protects technology infrastructure and introduces an entirely new category of cybersecurity risk.⁷⁰

However, as set out in the G7 Cyber Expert Group (CEG) statement on AI and cybersecurity, AI can also enhance cyber resilience. It is doing this through anomaly and fraud detection and response, predictive maintenance and patching and monitoring vendor and supply chain risk.⁷¹

Implications for the rest of the report

Most organisations need to build and maintain the strong cyber security foundations that will enable them to protect themselves against AI-enabled attacks. The FCA has established its system wide coordination approach in conjunction with industry with its Cyber Coordination Group (CCG) program.⁷² This focuses on getting the foundations right, sharing insights and building resilience. Coordination and collaboration on emerging threats within cyber risk will grow in significance as the threat from AI continues to escalate. The OECD suggests that the shift in cyber risk is moving to the level of financial stability concerns.⁷³

These shifts create pressure on how the UK's regulatory framework is applied in practice. The framework's principles remain sound, but its application becomes more complex as AI moves from assistive roles toward autonomous operation.

69 G7 Cyber Expert Group, [Statement on Artificial Intelligence and Cybersecurity](#), 2025.

70 Bain and Company, [Claude Mythos and the AI Cybersecurity Wake-Up Call](#), 2026.

71 G7 Cyber Expert Group, [Statement on Artificial Intelligence and Cybersecurity](#), 2025.

72 FCA, [Cyber Coordination Group insights](#), 2025.

73 OECD, [Cybersecurity and geopolitical risks in financial market](#), 2026.

Regulatory implications of AI-enabled finance

Regulatory implications of AI-enabled finance

Key findings

- Existing regimes remain effective for assistive AI, but some regimes come under strain as systems move toward human as Approver and Observer on our autonomy spectrum.
- Shared AI dependencies create ecosystem-level risks; concentration, correlated behaviours, and limited visibility that sit outside firm-level supervisory frameworks.
- Agentic supervisory model is essential both for firm based supervisory efficiency, and for monitoring and detecting system-wide risks.

The current framework: sound foundations under increasing strain

The UK's principles-based, outcomes-focused regulatory framework provides a credible foundation for an AI-enabled financial system.

Engagement with industry and other stakeholders indicates broad support for retaining this approach. The Consumer Duty, the Senior Managers Regime, operational resilience requirements and the wider conduct framework were designed to apply across evolving business models and continue to provide flexible tools for governing AI deployment within FCA regulated firms. Firms consistently emphasised that accountability should continue to rest with senior managers in regulated entities, exercised through systems, controls and oversight of AI-enabled decision making.

The question is therefore how the framework can be applied as the human role changes. In practice, its application becomes more complex as financial services evolve along the autonomy spectrum. At lower levels of autonomy, where AI supports human judgment, application of existing rules is relatively clear. As systems move towards greater delegation, with humans shifting from Operator to Approver or Observer, the application of the framework becomes more complex.

At higher levels of autonomy, firms may find it more difficult to evidence consumer understanding, demonstrate good outcomes, and maintain clear lines of accountability. Issues of consent, auditability, redress and the allocation of responsibility across delivery chains are likely to become more pronounced, where third-party models and platforms play a greater role. These developments do not undermine the framework itself but place increasing pressure on how it operates in practice.

The sections that follow explore how these pressures manifest across key regulatory regimes and where their application is most likely to be tested.

Key regulatory regimes: current fit and pressure points

The impact of AI on regulation is not uniform. Some regimes show early strain even at lower levels of autonomy, while others operate effectively until systems become more autonomous. The table below shows where each regime performs well and where pressures emerge as AI usage moves from assistive tool to autonomous operation.



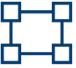







The analysis below considers each regime in turn.


Senior Managers Regime

The Senior Managers Regime ensures responsibility for regulated activity is clearly allocated to named individuals. Senior decision-makers are held accountable for how firms govern and oversee their businesses. Firms allocate responsibility to senior managers in key areas of their businesses, submit applications for FCA approval, and apply conduct rules that set out the enforceable standards for personal accountability.

The regime requires senior managers to take 'reasonable steps' to prevent regulatory breaches in their designated areas with responsibilities set out in a formal Statement of Responsibilities document that specifies their accountability.

Figure 16.
Where regulation comes under strain: the autonomy spectrum

	L1 Operator Human uses AI as a tool	L2 Collaborator Human and AI work together	L3 Consultant AI recommends, human decides	L4 Approver AI prepares, human authorises	L5 Observer AI acts within boundaries, human monitors
 Senior Managers Regime	Operates effectively Clear accountability through direct human oversight of AI-assisted tasks.	Operates effectively Accountability remains clear across human-AI workflows.	Operates effectively Accountability clear but may require effective human challenge of AI output.	Increasing complexity  Accountability remains clear but oversight of AI-enabled processes becomes more challenging.	Pressure point Meaningful human control likely to be difficult to evidence. Accountability remains clear but harder to exercise.
 Operational resilience	Increasing complexity  Early dependence on a small number of providers and infrastructure introduces shared risk. Firm level controls remain effective.	Increasing complexity Concentration risk increases as more firms rely on similar infrastructure. Most dependencies still identifiable at firm level.	Increasing complexity Shared models used for similar tasks may create correlated risk. Firm-level mitigation may become insufficient.	Pressure point Common points of failure likely to emerge. Disruptions may impact multiple firms simultaneously. Firm-level resilience may not be sufficient.	Potential system-level failure mode Risks likely to propagate rapidly across firms and markets. Coordinated response required but may be difficult at speed.
 Regulatory perimeter	Operates effectively AI use remains informational. Activity based framework aligns with discrete roles.	Increasing complexity  Influence emerges outside perimeter as AI shapes comparison, ranking and presentation of options. Still primarily informational but begins to affect choices.	Emerging Pressure Point AI provides personalised recommendations likely to resemble regulated activity but may sit outside the perimeter. Functional equivalence becomes visible.	Pressure point Activities influencing consumer financial decisions may occur outside the perimeter. Functional equivalence without regulatory oversight may emerge.	Potential regulatory framework mismatch AI platforms may become primary gateway of financial services. The perimeter may no longer capture firms with influence over consumer outcomes.
 Advice guidance boundary targeted support	Operates effectively Clear distinction between support that constitutes regulated financial advice or not.	Operates effectively Clear distinction between support that constitutes regulated financial advice or not.	Increasing complexity  Personalised recommendations may blur the boundary. Advice-like outputs may emerge at scale, often outside regulation.	Pressure point Continuous, adaptive recommendations may resemble regulated advice. Existing frameworks may need to adapt to accommodate AI capability.	Potential regulatory arbitrage Autonomous financial management may operate outside regulated financial advice. Boundary may no longer align to how decisions are made.
 The Consumer Duty	Operates effectively Good outcomes understood and measured can be evidenced using established methods.	Operates effectively Consumer outcomes remain demonstrable through existing approaches.	Operates effectively Personalisation raises questions on fair value and suitability across segments.	Increasing complexity  Consent and understanding become harder to evidence. Fair value and suitability more difficult to demonstrate.	Pressure point Continuous decision-making may make outcomes difficult to evidence. Consumer understanding uncertain. Existing consent model may be insufficient.

 First point of pressure for this regime

In a technological arms race scenario, the personal accountability regime could serve as an important mitigant against firms adopting a 'growth at all costs' strategy and internalising regulatory sanctions as a priced risk rather than a constraint. Under the Senior Managers Regime, individuals, not firms, remain personally accountable.

The Review agrees with most respondents that the Senior Managers Regime would be robust to the challenges of AI operating in the Operator, Collaborator and Consultant modes (Levels 1–3), where the regime continues to operate effectively. However, the Review acknowledges that without guidance, the increasing opacity inherent in aspects of the more delegated and autonomous operation of AI, human as Approver or Observer, coupled with factors such as model drift, could create an emerging pressure point in accountability.⁷⁴ That is, a potential gap between the outcomes of AI mediated decisions and the ability of the regulator to identify a de facto responsible party or for responsible individuals to exercise 'meaningful human control' where appropriate.⁷⁵

Assessments of the future adequacy of the Senior Managers Regime should account for two possible developments. First, the emergence of AI technologies mentioned in the systemic driver section could make future models less probabilistic and less opaque than current models. Second, assurance tools, including pre-deployment and ongoing checks, could better enable senior managers to take 'reasonable steps' to prevent regulatory breaches or misconduct, including risks from third-party models operating upstream.

Engagement paper responses and stakeholder feedback indicate that clearer guidance clarifying the reasonable steps required of senior managers could provide greater confidence in adopting delegated and autonomous AI. Mass adoption could help to make markets work well and support long term growth and competitiveness.

Operational resilience

The Operational Resilience framework requires firms to identify their important business services (services that, if disrupted, could cause intolerable harm to consumers or threaten market integrity) and set impact tolerances for how long these services can be disrupted before harm occurs.

The Critical Third Party (CTP) regime aims to reduce systemic risk and strengthen the operational resilience of the financial sector by bringing in the most significant third party service providers under direct regulatory oversight. HM Treasury can designate these providers whose failure or disruption of services to firms could threaten stability of, or confidence in, the UK financial system. Once designated, the regulators (the FCA, PRA and the Bank of England) oversee the systemic third-party services those providers deliver to the sector.

⁷⁴ Santoni de Sio, F., Mecacci, G. Four Responsibility Gaps with Artificial Intelligence: Why they Matter and How to Address them. *Philos. Technol.* 34, 1057–1084 (2021). <https://link.springer.com/article/10.1007/s13347-021-00450-x>

⁷⁵ Methnani, L., et al, Let Me Take Over: Variable Autonomy for Meaningful Human Control, Vol. 4, 2021 | <https://doi.org/10.3389/frai.2021.737072>

The operational resilience framework establishes firm-level accountability for resilience, while the CTP regime extends regulatory oversight in a way that complements, but does not replace, firms' responsibilities to manage the potential impact to their firm of the failure or disruption of a third-party.

Firms can map AI dependencies on models and cloud infrastructure, set impact tolerances, maintain contingency arrangements, and test resilience through scenario testing. The CTP regime is technology agnostic, so it can capture major cloud providers and providers of services based on new, emerging or rapidly developing technologies, where designation criteria are met.

AI may introduce new potential systemic risks that sit outside the scope of firm-level resilience frameworks. Three specific risks may emerge:

First, shared model dependencies may create correlated, system-wide risks. When multiple firms rely on the same AI model provider, an outage, security breach or model degradation could affect all reliant firms simultaneously.

Second, correlated behaviours may amplify market risks. Where firms use similar models for credit underwriting, investment recommendations or pricing, synchronised decision-making could amplify market movements, create procyclical effects, or lead to coordinated withdrawal of services from particular customer segments.

Third, regulators may have limited visibility of concentration in model provision and infrastructure across the financial system, making it more difficult to identify vulnerabilities or coordinate responses quickly until disruptions occur.

These are ecosystem-level risks, not firm-level risks. Existing regimes continue to function as designed. The issue is that AI may introduce a new category of systemic risk through shared dependencies that could require greater coordination across firms and regulators particularly, at higher levels of autonomy where common points of failure and system-wide propagation become more likely.

Regulatory perimeter

The regulatory perimeter defines which financial activities require FCA authorisation. Firms come within the perimeter if they engage in a specified activity in relation to a specified product or investment and do so 'by way of business,' that is, continuously and for commercial purposes. HM Treasury defines the perimeter; the FCA can recommend changes based on monitoring and evidence.

The activity-based foundation remains broadly sound. It captures regulated activities regardless of the technology used. Firms using AI to provide advice, arrange transactions or promote financial products are subject to the same tests as firms using other delivery methods.

AI creates increasing complexity and emerging pressure points that the activity-based perimeter was not designed to address:

1. Influence over consumer outcomes may increasingly emerge from outside the regulatory perimeter, as general-purpose AI applications shape consumer understanding and decisions without clearly undertaking regulated activities, creating a potential gap between where influence sits and where protections apply if something goes wrong.
2. Competitive asymmetry may emerge where regulated firms face obligations under conduct rules and financial promotion requirements, while platforms and model providers may exert similar influence without equivalent obligations, potentially creating an uneven playing field.
3. Commercial models may increasingly blur boundaries. As subscription tiers, referral arrangements, paid placement and preferential routing develop, functionally similar activity may become established outside the perimeter before oversight adapts. As interactions become continuous and blended within AI-mediated journeys, activity-based tests designed for discrete roles may become less effective.

AI platforms are expanding consumer access to financial information, enabling consumers to understand products, compare options and navigate financial decisions more easily. This creates new routes to market and improves access for consumers who may not engage with traditional channels.

While the Review has not observed technology firms seeking to enter financial services as a core business line, the evidence from the engagement survey indicates that consumer reliance on AI platforms for financial information is material and growing, and that regulatory boundaries may require further clarification to maintain effectiveness while enabling beneficial innovation to continue.

Consumer protection law may in some cases provide redress for AI-related harms beyond our perimeter. However, the Law Commission's 2025 Artificial Intelligence and the Law report found that while consumer protection legislation is technology-neutral, its effectiveness can be context-dependent and may require further scrutiny and reform.⁷⁶

Advice guidance boundary targeted support

The advice/guidance boundary defines what constitutes regulated financial advice versus guidance. Regulated advice involves personal recommendations based on individual circumstances and requires FCA authorisation. By contrast, financial guidance more often involves providing information or narrowing options without going as far as making personalised recommendations.

Targeted support, introduced following the Advice Guidance Boundary Review, enables regulated firms to make recommendations designed for groups of consumers with common characteristics to help them make important decisions around their pensions and investments.

⁷⁶ Artificial Intelligence and the Law <https://lawcom.gov.uk/news/artificial-intelligence-and-the-law-a-discussion-paper/>

This should increase access beyond the approximately 9% of consumers who receive traditional regulated financial advice, though recommendations remain generalised rather than hyper personalised.⁷⁷

In respect of discrete human interactions, the boundary operates clearly. Targeted support provides a structured way for firms to offer enhanced support within defined parameters.

However, AI general-purpose LLMs and frontier models have the potential to expand access even further by providing hyper personalised recommendations based on individual characteristics and data, addressing long-standing advice gaps for consumers who cannot access traditional advice or benefit from pooled targeted support. This could be a positive development for financial inclusion.

However, increasing complexity and emerging pressure points arise as AI capabilities expand:

- **Personalised outputs challenge the boundary.** General-purpose LLMs and frontier models may provide hyper personalised recommendations sourced directly from individual consumers, delivering advice-like support that goes beyond what targeted support enables while operating outside the regulatory perimeter.
- **Regulatory arbitrage.** Regulated firms report they are restricted from providing similar personalised support while unregulated AI platforms may exert comparable influence without equivalent obligations, potentially creating competitive asymmetry and consumer protection gaps.
- **Future interfaces blur responsibility.** As covered in the consumer journeys and the competition chapters, consumers are likely to access support through intermediaries, devices or AI interfaces. When they do, it may become less clear where responsibility sits when advice-like services are delivered with increasing personalisation through continuous AI-mediated interactions.

In the recommendations section, we stress the need to monitor the advice guidance boundary in the future, as part of the perimeter review, but also later if AI increases in capability and risks of error are reduced, to rethink whether or not targeted support or other measures can be adjusted to better enable firms to use AI on a more individualised basis. This would have benefits for tackling the advice gap and therefore, unlocking the microeconomic benefit of millions of more informed financial decisions being made by UK citizens, while noting that at higher levels of autonomy the risk of regulatory arbitrage may increase where activity no longer aligns with the existing boundary.

The Consumer Duty

The Consumer Duty requires firms to act to deliver good outcomes for retail customers. It establishes that firms must provide products and services that are designed to meet customers' needs and that they know provide fair value.

⁷⁷ [Financial Lives 2024: Key findings from the FCA's Financial Lives May 2024 survey](#)

Firms should also communicate and engage with customers so that they can make effective, timely and properly informed decisions and can take responsibility for their actions and decisions. Under the consumer support outcome, firms should enable consumers to realise the benefits of the products and services they buy.⁷⁸

The Duty applies straightforwardly where AI supports human decisions. Firms can more easily demonstrate good outcomes, consumer understanding and foreseeable harm when AI systems operate in Operator, Collaborator and Consultant modes. As consumer journeys become more dynamic, personalised and delegated, demonstrating compliance becomes more complex. Increasing complexity and emerging pressure points arise as autonomy increases:

On price and value, AI-enabled pricing may make it harder to distinguish benign personalisation measures from extraction where consumers pay more without receiving additional value.

- On consumer support, AI could improve availability, but firms may still need human involvement for complex or sensitive issues.
- On consumer understanding, dynamic and personalised journeys may make it harder to evidence that consumers genuinely understand products, decisions and risks.
- On products and services, AI increasingly shapes design and delivery, which may make it harder to ensure outcomes remain appropriate and fair across different consumer groups, particularly where models may embed historic patterns of bias.

At higher levels of autonomy, these challenges may intensify. Consumer understanding may become uncertain where AI makes thousands of micro-decisions over time. One-off consent may prove insufficient for ongoing autonomous action, and detriment may occur before consumers are aware.

Engagement paper responses support the continued relevance of the Consumer Duty as the framework for AI-mediated financial services. Stakeholders highlighted that firms need clearer interpretation of how to evidence outcomes in dynamic journeys and how to ensure consent remains meaningful where systems act continuously, particularly at higher levels of autonomy where demonstrating outcomes and understanding may become more difficult.

Enabling infrastructure for agentic finance

Throughout this report, we have discussed AI assistants, AI agents and AI interfaces taking on ever more activity within firms and for consumers. However, we have also noted the concerns raised by consumers around trust, control, access and liability - what happens when things go wrong? This section is about establishing trusted protocols to help support safe and trusted, innovative and growth-enhancing agentic finance.

⁷⁸ Our expectations of firms under the Duty, See 1.9 in [FG22/5: Final non-Handbook Guidance for firms on the Consumer Duty](#)

We have observed that many payments providers are seeking to solve these questions for 'agentic payments'.

Agentic payments are payments initiated or completed by an AI agent on someone's behalf. For example, you might ask an AI agent to 'find me a train journey from London to Cardiff this evening and book me a standard fare'. The agent might look at your diary, assess your ability to get to Paddington Station and make the payment for the ticket using permissions you have granted.

In a small business, an agent might pay an invoice or settle bills. Agentic payments can use cards, bank transfers, wallets or even stablecoins. Payments firms are developing agent tokens, authentication, spending controls and records that can trace and identify the agent, the consent and the transaction. Agentic payments is the point at which the agent is authorised to execute and complete a transaction. The UK government recently announced that agentic payments form part of key measures to improve growth and innovation in the economy and in financial services.

Agentic finance is wider than agentic payments. It is a broader concept that encompasses agentic payments but includes AI agents who will help consumers manage their financial lives. An agent might monitor deposits and savings, identify better insurance products, manage investment portfolios, support lending decisions and manage debt repayments, make, or manage, claims or complaints and switch providers. This agent is more comprehensive and covers a wider range of tasks including the search, info gathering, recommendation and taking of action across financial services, not just executing payments.

At the heart of agentic finance is enabling infrastructure that needs to support safe and accountable delegation to the AI agent.

This enabling infrastructure matters hugely for consumer protection, for growth and innovation and for the contribution of agentic finance to the wider economy. It could unlock significant productivity gains by reducing frictions in the system, streamlining operations and reinforcing the UK's position as a global fintech hub.

Some commentators identify autonomous agents as a major source of future economic growth but explain that they will require access to capital and financial services. From a consumer perspective, agentic finance could deliver a host of benefits including addressing longstanding issues such as the advice gap and increasing access for underserved consumers.⁷⁹

The table on the next page sets out the building blocks for agentic finance, and it recommends that frameworks for data, identity, authority to act, execution and liability need to operate together if systems are to move beyond assistive use, that is, human as Operator and Collaborator. Without this, adoption will remain partial, fragmented or reliant on repeated human intervention, and the economic and consumer benefits will be curtailed or delayed.

⁷⁹ <https://www.weforum.org/stories/2024/12/agentic-ai-financial-services-autonomy-efficiency-and-inclusion/>

The regulator and the market will need to consider how these foundations (discussed in the agentic finance infrastructure annex) will evolve to shape the development of safe agentic financial services.

We recommend in our priority recommendation 5 that the FCA works to develop a trusted agent protocol to enable an agentic finance ecosystem to develop in a safe and trusted way that supports consumers and realising significant economic opportunity for firms and UK growth and productivity.

The building blocks of agentic finance

Building block	Implication for agentic finance
Data	Agents depend on goal, observation and outcome data. Data quality and availability remain a binding constraint on scale.
Identity	Agents are linked to a consumer and mandate, with verifiable authority over actions taken.
Authorisation and delegation	Decision-making shifts from point-in-time approval to pre-authorised, bounded and revocable mandates operating over time.
Execution (payments)	Existing payment frameworks assume human approval and are not aligned to delegated or autonomous agent execution.
Liability	Responsibility becomes harder to attribute as actions are delegated, shaping behaviour, trust and adoption.
Supervision and audit	Agent activity becomes traceable, with visibility of authority, actions and outcomes supporting redress and system-wide oversight.

Implications for the regulator and supervisory model

We have looked at how the regulatory framework might need to change in the future, but a more pressing area is change within the regulator itself. The agentic supervisory model is recommended in the next chapter because AI is likely to become a core operational capability for regulated firms and it will transform how consumers engage with financial services.

This agentic supervisory model is needed for three strategic reasons. The first is efficiency. AI can make regulation more effective and efficient across authorisation, supervision and enforcement. Faster application processing, real-time monitoring, and automated routine processes enable the FCA to deliver on the secondary objective for international competitiveness and growth and also to underpin protection of consumers through supporting firms to achieve better outcomes for their customers more efficiently.

Second, the capability to tackle more harm and enable more growth. AI-enabled tools allow the FCA to detect risks earlier, intervene faster, and process more cases. This means preventing harm at scale.

The third is keeping pace with technological change. Across regulated firms, AI will become a core operational capability. This means the speed, pace and capability of firms will improve, productivity will improve as we have seen AI capability is a system driver for firm adoption and improved productivity, innovation and output. Likewise, consumers will go on their own journey subject to the control, trust and access points we have made.

We also see from system shift 4 that external actors in fraud, financial crime and cyber will move faster, often spurred by new AI model capabilities. And we know that multi-agentic AI systems will proliferate and these systems can lead to a multitude of system-wide risks that will either amplify existing risks or bring new risks but with greater speed, scale and capability.

We are likely to see a shift from human episodic, often retrospective, intervention to a more dynamic, system-wide, near continuous real-time monitoring and more proactive intervention, supervisory engagement and enforcement. We consider it will be important that an AI agentic supervisory model will need to be developed to continue to fulfil the FCA's objectives in future.

In other words, the regulator itself needs to move down the same autonomy spectrum as firms, consumers and the wider economy towards AI as Collaborator and AI as Approver.

An AI agentic supervisory model would operate continuously across the FCA's remit, monitoring, flagging, triaging and surfacing firm and system-wide signals. This evolution has two components: enhancing firm-level supervisory efficiency and building capability for new risks and system-wide oversight.

Enhancing firm-based supervisory efficiency

Firm-by-firm supervision remains the foundation of the regulatory model. Firms are responsible for the services they provide and the outcomes they deliver, including those that involve AI.

But as firms use AI to process more data and make more decisions, the nature of supervisory engagement changes. Supervision needs to be timelier, and the FCA needs to understand how firms design, monitor and govern AI-enabled systems as they evolve.

The core principle does not change. Human supervisors should continue to be responsible for challenging firms and making decisions on intervention. The FCA's own use of AI remains subject to the same expectations it puts on firms, including explainability, auditability and appropriate human oversight.

Building capability for system-wide risk oversight

As AI systems move towards autonomous action across financial services, more material system-level risks will start to emerge, testing the limits of today's firm-by-firm supervisory model.

These risks are driven by shared factors across firms such as concentration of resources, increasing automation, greater interconnectedness, and evolving complexity.

The reliance on a small number of concentrated AI providers, models, and infrastructure. Firms' responses to our Engagement Paper noted dependency on a core group of US-based models and a handful of hyperscalers. This raises clear resilience concerns, as the failure of a single provider could disrupt many firms at once. At system scale, a subtler risk emerges, as firms relying on the same models may behave in similar ways, leading to herding and a narrowing of consumer outcomes.

As automation develops, flowing through increasingly connected services, harms can propagate more easily. With many firms intending to expand their use of agentic AI, these issues may spread quickly to affect multiple firms and large numbers of consumers before they are detected.

Increasing complexity may also make risks harder to observe. AI systems can be difficult to interpret, which can obscure bias and weaken oversight, raising conduct risk where poor outcomes go unseen. Over-optimisation for performance is likely to produce unintended outcomes at scale, particularly as decisions become hyper-personalised and less comparable across consumers.

Each of these dimensions turns a firm-level problem into a system-level one in that concentration aligns firms' exposures, automation lets harms travel between them, and complexity hides the build-up. None of it is visible firm by firm, so supervision will need to look across the system as a whole, monitoring correlations between firms, mapping shared dependencies, and stress-testing how systems behave together.

Data quality: foundational for firms and regulators

Data quality is foundational for AI-enabled supervision. Effective system-wide oversight depends on firms providing timely, structured, high-quality data that enables cross-firm pattern detection and near real-time monitoring. The FCA's own use of AI across authorisation, supervision and enforcement also depends on data quality, not just data availability.

Detailed analysis of data quality requirements for firms, consumers and regulators, including the system shift towards trusted, traceable data in regulated financial services, is set out in the annex.

International coordination: a systemic imperative

AI models, cloud infrastructure and platforms operate across jurisdictions, creating dependencies that no single regulator can address alone. If a major model provider experiences a security breach, widely used models generate biased outputs, or shared infrastructure fails, impacts can propagate across jurisdictions simultaneously.

Regulators will therefore need to share information earlier and coordinate collective responses when AI-related risks arise across borders.

We spoke to regulators across multiple jurisdictions, governmental departments, international standard setting bodies, NGOs and academics. While stakeholders recognised the potential benefits of AI, they also recognised the potential risks, some of which are cross-border. These include those that arise from shared dependencies on global providers, AI-enabled financial crime and cross-border data flows. During our engagement, there was a universal willingness to share experiences, coordinate on standards and build mechanisms to respond to cross-border incidents.

The OECD AI principles (first adopted in 2019), later developed into more granular guidance and reporting mechanisms via the Hiroshima AI process, are an important foundation for international convergence and have become an important global reference point for national AI strategies.^{80 81}

But despite broad agreement on principles, international coordination on AI governance remains a complex challenge. Jurisdictions are diverging in their regulatory approaches, have different levels of national AI capability and have varying national priorities around consumer protection, industrial strategy, national security and technological sovereignty. International financial standard-setting bodies are treating AI as a material issue, monitoring developments closely and producing initial resources - such as IOSCO's recent supervisory toolkit and the FSB's guidance on sound practices for AI adoption - but are not yet seeking to set comprehensive standards.^{82 83}

In the absence of fulsome national or international frameworks, firms often look to voluntary standards such as those developed by the National Institute of Standards and Technology (NIST), the International Organisation for Standardisation (ISO) and other bodies.

Over time, there will be a need for more structured international supervisory co-ordination and, where appropriate, exploration of sector-specific international standards to complement national frameworks.

Moreover, effective monitoring and mitigation of system-level risks will require closer dialogue with key AI ecosystem actors. Recent engagement between standard-setters and model providers around controlled access to frontier models (such as the Project Glasswing case) suggests that there may be scope for closer dialogue between financial authorities and AI providers where such risks emerge.

Closer international co-ordination should also promote alignment on specific frameworks for agentic protocols to enable responsible cross-border agentic activity.

Further detail on emerging international regulatory approaches and coordination considerations, as well as varying national approaches to AI adoption, is set out in the international approaches annex.

80 AI Principles, OECD, <https://www.oecd.org/en/topics/sub-issues/ai-principles.html>

81 Hiroshima AI Process, <https://www.soumu.go.jp/hiroshimaaiprocess/en/index.html>

82 IOSCO, [IOSCO publishes AI Supervisory Toolkit 2026](#)

83 FSB, [Sound Practices for Responsible Adoption of Artificial Intelligence \(AI\): Consultation report 2026](#)

Conclusion: progressive adaptation, not wholesale replacement

The current regulatory framework does not need to be replaced. Its foundations remain sound where AI supports human decisions. The framework was designed for human-intermediated financial services where decisions made by people, approved by people, with clear lines of accountability to named individuals. AI changes this.

As autonomy develops, the way the framework applies is not uniform. In some regimes, particularly the Senior Managers Regime and Consumer Duty, existing arrangements continue to operate effectively where human control remains central.

In others, including operational resilience and the regulatory perimeter, pressure is already evident at earlier stages as AI reshapes dependencies, market structure and influence over consumer decisions. A key inflection point emerges as AI moves from supporting decisions to shaping them. At this stage, complexity increases and existing boundaries begin to blur. As firms move further towards delegation and autonomous operation, this develops into material strain, and in some cases structural mismatch.

As this transition occurs, the framework comes under increasing strain. The FCA, as regulator must balance its risk appetite for the potential opportunities AI might bring to consumers, competition and innovation, growth and productivity against some of the risks to market integrity, to consumers, in particular those with vulnerable characteristics, financial capability or access issues, and how best to mitigate those.

Our recommendations focus on underpinning the regulatory framework to support adoption of AI in a safe and trusted way while supporting growth and innovation.

Accountability might become harder to trace across interconnected AI systems and third-party model providers. Influence over consumer outcomes sits increasingly outside the regulatory perimeter. The FCA and sectoral regulators may lack tools to tackle system-wide harm that propagates across firms. Consumer protection standards designed for discrete, human-intermediated interactions may not hold in continuous, AI-mediated journeys. And international coordination will become ever more important.

Navigating toward the opportunities of agentic finance while managing these risks requires specific foundations to be put in place:

- Systemic oversight of model providers and shared infrastructure, enabling earlier detection of concentration risk, correlated behaviours and cross-firm harm
- Clarity on the regulatory perimeter, ensuring that influence over financial outcomes faces appropriate oversight regardless of where it sits
- Consideration of additional powers for the FCA and sectoral regulators to address horizontal, system-wide harm that current firm-by-firm supervision cannot fully capture

- Adaptive guidance on the Consumer Duty and senior managers regime, clarifying how accountability is maintained and consumer protection standards apply as delegation increases
- Measures to protect access and capability for vulnerable consumers, ensuring AI-enabled services do not widen existing inequalities.

Signals will need to be monitored closely. Where systems move beyond level 4 toward full autonomy (level 5), certain frameworks will require adaptation. The pace of that transition is uncertain. The direction is clear. The regulatory response must be evidence-led, proportionate and ready to evolve as adoption progresses.

These findings show that the regulatory framework remains a strong foundation, but it must adapt in a disciplined and progressive way as AI becomes more embedded across financial services. This creates an opportunity for the UK to support safe adoption, enable innovation, and strengthen its position as a leading and trusted financial centre.

The recommendations that follow are the practical response. They maintain and adapt the existing framework so it can continue to deliver good outcomes for consumers while supporting innovation, competition and growth.

Priority recommendations for AI-enabled finance

Priority recommendations for AI-enabled finance

From analysis to action

The analysis points to clear opportunities to support greater productivity and growth in UK financial services. The FCA should proactively adapt the regulatory framework to enable safe, trusted and innovative AI adoption by firms and consumers, while mitigating emerging risks.

This growth will largely come from firm productivity, efficiency and innovation but we also consider that supporting more consumers to make better financial decisions would also contribute to economic growth.

Underpinning our recommendations is our conclusion that the regulatory framework remains fit for purpose. We also consider that the FCA's principles-based and outcomes-focused approach - as exemplified by the Consumer Duty, and its groundbreaking innovation work in the AI Lab - will enable the FCA to proactively adapt to AI to 2030 and beyond.

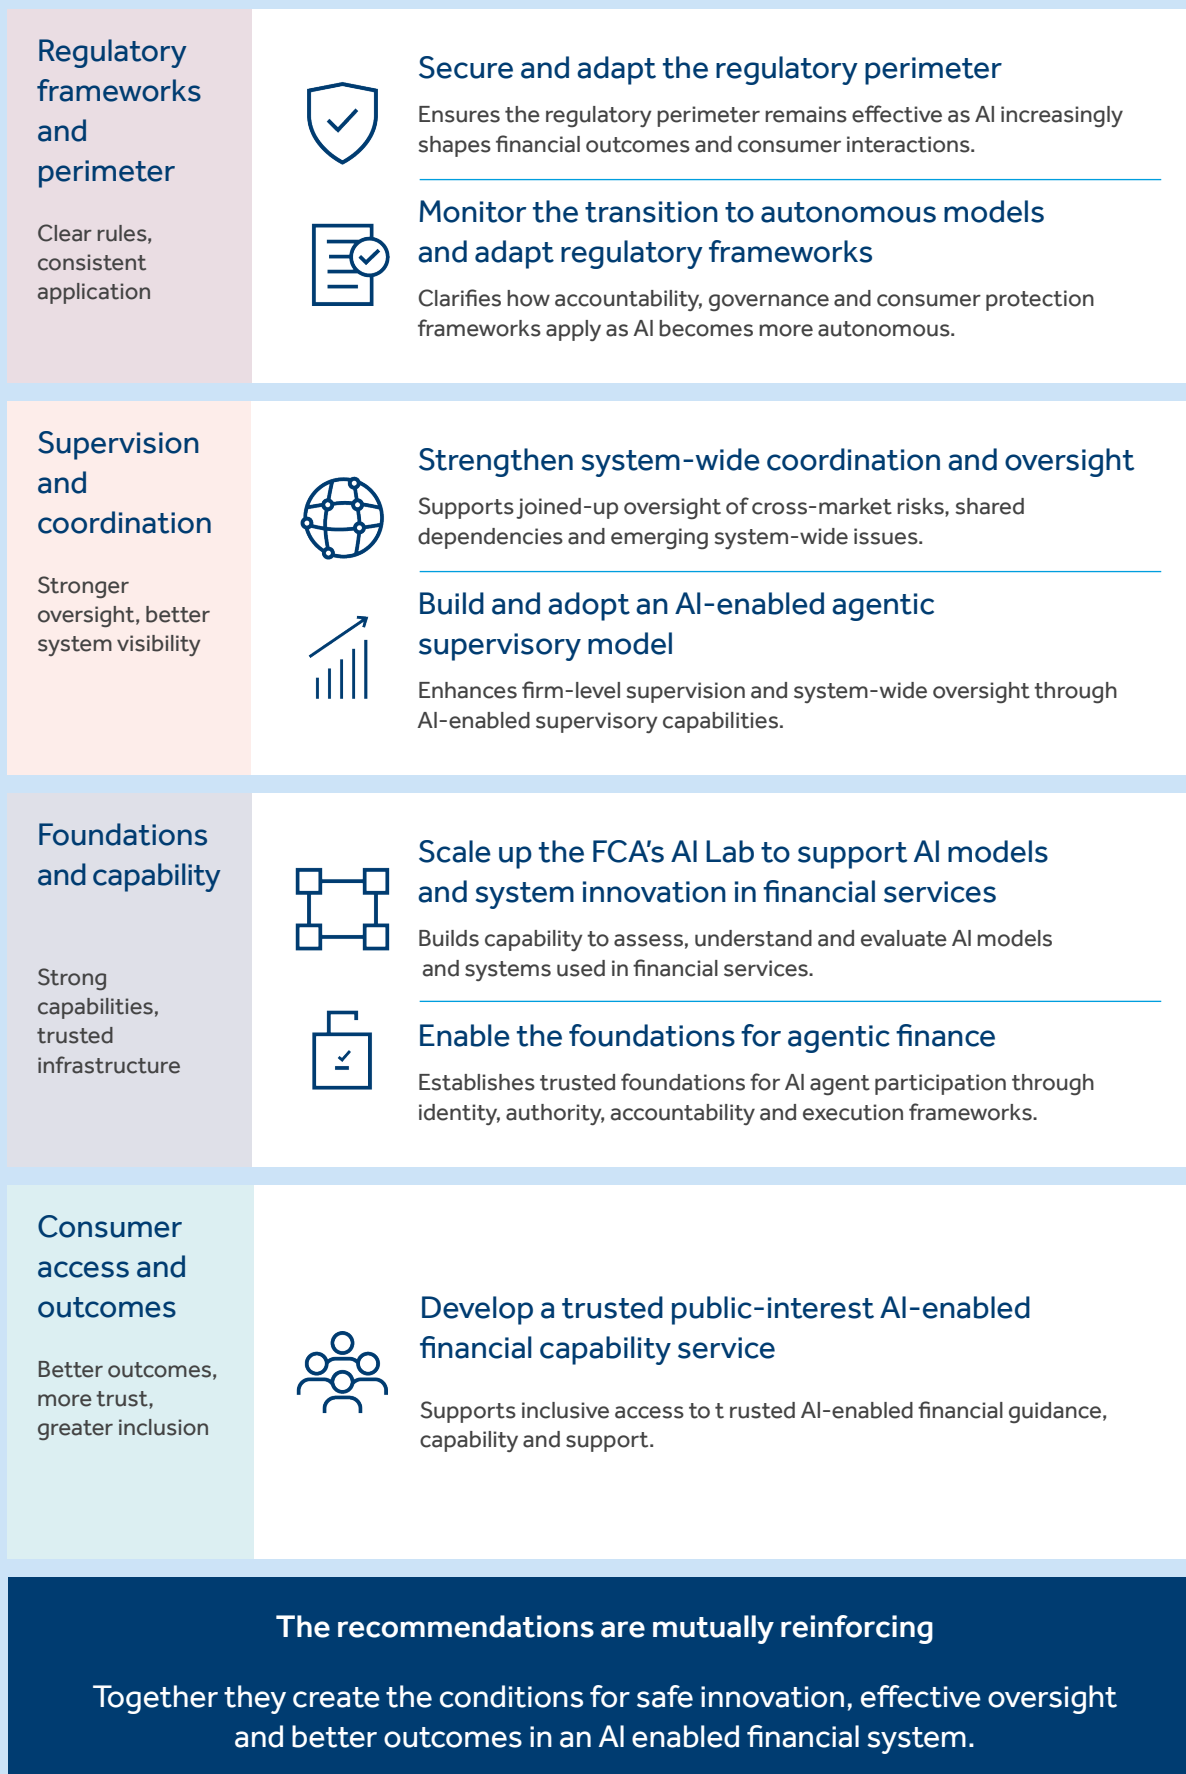
But we see strategic choices and risks for the FCA because, on our autonomy spectrum, existing regulatory levers continue to operate effectively at the Operator, Collaborator and Consultant stages as AI supports, assists and collaborates with firms and consumers. Complexity and pressure points emerge as AI moves to the Approver or Observer stage - that is, when it becomes a core operational capability embedded in more workflows, processes, and decision-making, as well as across financial services.

This chapter sets out seven priority recommendations for the FCA Board to consider with the Executive. The recommendations support the FCA's statutory objectives by aiming to maintain effective consumer protection, preserve market integrity and promote competition, and international competitiveness and growth, as AI transforms financial services.

We recognise that AI is changing rapidly and that its adoption in financial services is also moving at pace for firms and consumers. The recommendations can therefore be adapted to meet capacity, risk appetite and changing external factors, but in essence, they cover four interconnected areas: Regulatory frameworks and perimeter; Supervision and coordination; Foundations and capability and consumer access and outcomes.

Action across these four broad areas, we consider, is required to balance the risk and opportunity AI presents for financial services.

Figure 17.
How the recommendations operate together across an AI-enabled financial system



Priority recommendation 1: secure and adapt the regulatory perimeter

Recommendation

Short-term recommendation: The Mills Review recommends that the FCA should consider securing and adapting the regulatory perimeter for AI-mediated retail financial services by launching a review into the scale, nature and impact of general-purpose LLMs outside the perimeter. This is so that the regulatory boundary continues to reflect how financial services are delivered and where influence over financial services outcomes is exercised in practice.

The FCA should consider conducting a review within three to six months of the date of this report. The review should examine how consumers use AI including general-purpose LLM tools for personal financial management across savings, investments, pensions, mortgages and debt management, and the implications for competition, innovation and growth. It should examine the risks of consumer harm, including how far its usage has or will move along the autonomy spectrum, and any impacts on market integrity and the potential for regulatory arbitrage.

Bringing together policy and competition experts, it should also include legal expertise to examine how current legal and Handbook frameworks, including Perimeter Guidance (PERG) apply, where gaps emerge. This includes how advice and guidance boundaries operate in conversational interfaces, how by way of business tests apply to general-purpose tools, and how financial promotions and arranging rules apply within continuous AI-enabled journeys.

Based on this evidence and analysis, the FCA should determine whether to amend guidance, recommend perimeter changes to government, or maintain the current approach.

Long-term recommendation: The FCA should keep under review the impact of frontier model capabilities and increasing AI adoption on the effectiveness of the activity-based regulatory perimeter. The UK regulatory regime is based on activities: Firms obtain permissions from the FCA to carry out certain activities and the FCA can regulate those activities regardless of whether the entity carrying them out is a bank, an insurer, a technology firm or even a dentist.

But, the pace of adoption and improvements in AI capability could, in some scenarios, lead to greater pressure on the activity-based regime and gaps in regulatory frameworks at the FCA or cross-sectoral regulators.

AI brings system-wide, cross-sectoral risks and opportunities and general-purpose technologies can bypass the rules by potentially providing economically equivalent products and services which fall, due to an activity-based regime, outside the perimeter. These pressures are likely to continue to emerge as use and adoption of agentic AI increases, especially once AI agents can take some action on behalf of consumers.

To mitigate this risk, we also recommend that the FCA consider requesting the government boost the FCA's existing powers under the CTP regime (currently restricted to systemic issues) and the DAR to meet the challenges mentioned above, and giving direct powers under the DMCC to the FCA and other sectoral regulators.

Why this matters

AI increasingly shapes how consumers understand products, compare options and make financial decisions through conversational interactions that resemble guidance or advice. This affects consumer outcomes and consideration of who might be responsible for them.

While AI can expand consumers' access to financial information and create new routes to market, particularly for those who may not ordinarily engage with traditional advice channels, it can simultaneously extend the influence of firms undertaking financial actions beyond the current perimeter.

Interactions increasingly take place through general-purpose tools such as ChatGPT, Gemini and Claude, shaping how consumers interpret information, narrow options and act. Our survey evidence suggests a potential gap between consumer expectations and regulatory reality: many consumers do not recognise that they may have limited or no formal recourse to redress if something goes wrong, while others assume protections exist where they do not.

In the investments sector, targeted support permissions allow firms to provide recommendations to groups of consumers with shared characteristics. This differs from outputs generated by general-purpose AI tools, which may be based on individual prompts, data and broader knowledge of the consumer.

However, it should be noted that advice/guidance boundary issues mentioned above in respect of investments, apply equally in a number of other retail finance areas: debt management, savings, pensions and drawdowns, mortgages and home finance.

Evidence from the Review's consumer research shows that consumers are already using general-purpose AI tools across these areas to support financial decisions, often in ways that resemble guidance or advice. So, it is important to note that the consumer protection issues could run across many sectors even where competition effects differ.⁸⁴

As AI becomes more embedded in consumer journeys and commercial models evolve - through subscription tiers, referral arrangements, paid placement or preferential routing - functionally similar activity to what we recognise as regulated activity may become established outside the perimeter before oversight adapts.

⁸⁴ Relevant FCA Handbook provisions include, for example: COBS 9 and 9A (suitability), and COBS 10/10A (appropriateness in non-advised investment services), MCOB 4 and 4A in relation to mortgage advice and sales, CONC 8 in relation to debt advice and debt counselling, relevant provisions on financial promotions (COBS 4 / MCOB 3A), and the advice/guidance boundary as set out in the FCA's Perimeter Guidance Manual (PERG), in particular PERG 8.

This creates practical issues about when AI-mediated interactions move from mere unregulated information provision, into activity that resembles regulated advice, arranging or promotion, and how consistently those boundaries can be applied.

There is a question to address about whether the activity-based framework continues to capture where financial influence is exercised in practice.

Key actions

The Mills Review recommends that the FCA should consider adopting a staged approach to active perimeter management.

The foundational action and our priority recommendation is review of the perimeter, and we recommend this is finalised within the next three to six months.

1. Ongoing monitoring: The FCA should also consider establishing ongoing monitoring of consumer reliance on AI outputs, evidence of harm, movement along the autonomy spectrum, and evolving commercial models.
2. Market engagement: Engage proactively with platforms, providers and firms to understand new business models and identify where activity may move closer to the regulatory perimeter. Engagement should be evidence-led and should enable the early detection of risks, so as to reduce the likelihood of disruptive intervention.
3. Intervention: Prepare graduated intervention tools as risks emerge, using existing mechanisms to target specific activities without requiring full authorisation.⁸⁵ Establish escalation pathways from monitoring through to supervision, intervention or enforcement. Thresholds should be evidence-based, focusing on observed consumer harm, market distortion or regulatory arbitrage.
4. Strengthening coordination and MOUs: to support cross-regulator and cross-government oversight of AI activity outside the perimeter. The FCA should strengthen MOUs and cooperation mechanism with other agencies who have ex-perimeter powers. These include strengthening coordination with HMT in relation to the DAR and in respect of the Digital Markets, Competition and Consumer Act, where the CMA has powers to act.
5. New powers: Then, as above, our long-term recommendation to advocate for new powers for the FCA (direct DAR powers) and the FCA and other sectoral regulators (direct DMCC powers).

Implementation considerations

The review should commence as soon as practicable. Consumer reliance on AI is growing, commercial models are evolving rapidly, and early clarity is needed to support responsible innovation while protecting consumers. This clarity can be delivered through guidance without legislative change.

⁸⁵ This includes the Designated Activities Regime, where HM Treasury can bring specific activities within the regulatory perimeter on the FCA's recommendation.

Managing the perimeter can be done within existing FCA powers, though effectiveness depends on coordination with the CMA on platform conduct, engagement with government on legislative clarification, and aligning with international bodies through IOSCO, FSB and bilateral engagement to share evidence, reduce compliance complexity and limit regulatory arbitrage.

Priority recommendation 2: strengthen system-wide coordination and oversight

Recommendation

The Mills Review recommends that the FCA consider enhancing and establishing coordination across domestic authorities and international partners to address the risks and opportunities that AI could deliver.

Establishing safeguards for delivering innovation and growth opportunities is required, but in the absence of a dedicated AI regulator or a dedicated AI cross-sector regime including legislation, coordination is an imperative for UK sectoral regulators. It should take place around resilience, data, competition, security, safety risks, consumer protection, supervisory approaches and standards-setting.

It might require deeper international cooperation on cross-border AI dependencies, collective monitoring of concentration and ecosystem vulnerabilities, including those arising from highly concentrated upstream providers of models, cloud infrastructure and distribution platforms, and a coordinated response to incidents affecting multiple firms or jurisdictions simultaneously.

The FCA should consider monitoring whether the current coordination frameworks remain fit for purpose and, if appropriate, recommend to government options to rationalise and strengthen cross-regulator coordination. The FCA should also consider whether new coordination mechanisms or forums are required where existing arrangements are insufficient to address emerging AI-related risks.

Why this matters

AI concentrates some risks at system level through shared models, infrastructure and cross-border dependencies. If a major model provider experiences an outage or security breach, multiple UK financial institutions relying on that provider could be affected at the same time.

No individual firm may be able to fully mitigate that risk through firm-level controls alone. The issue sits at ecosystem level and would require coordinated monitoring, shared understanding, and a joined-up response.

This also links to the Review's analysis of the Operational Resilience framework and the Critical Third Party (CTP) regime. As AI adoption increases, concentrations in model providers, cloud infrastructure and other upstream technology services may become more important drivers of system-wide resilience risk, raising questions about how regulators monitor and understand shared dependencies across the financial services ecosystem.

Without stronger coordination, there may be an increased likelihood of common points of failure. Harms can propagate rapidly across firms through correlated model behaviour, shared infrastructure failures, or new fraud and conduct risks that scale faster than individual regulators are able to detect. Cross-border disruptions may reach UK markets without adequate warning or response.

Key actions

1. **Strengthen domestic coordination:** We identify three categories where AI-related system-wide risks may emerge and coordination should strengthen across UK authorities through more formal use of existing mechanisms and engagement with government departments such as DSIT and authorities such as AISI.

For each of these categories, we recommend that the FCA seeks, with partners, to identify the role of relevant forums, committees and networks in relation to AI system-wide risks to ensure that there are no gaps in coverage.

The FCA and its cooperation partners across each of the three domestic areas should consider whether there is scope for rationalisation of agendas and activities to ensure maximal use of capability and resource.

- a. **Financial services-centred risks:** Continue with the strong coordination on AI as shown through the AI Consortium convened by the Bank of England and the FCA with industry; coordination and cooperation across regimes such as Operational Resilience and Critical Third Parties; senior leaders attending committees such as the FCA's CEO at the Financial Policy Committee and the CEO of PRA attending the FCA Board.

Together with the Regulatory Initiatives Grid and other coordination mechanisms, we consider that the risks posed by AI to financial services can continue to be managed through firm-level supervision and regular cross-authority coordination.

The FCA and PRA might wish to consider more formal mechanisms which bring together risks across the sector holistically, particularly where they affect multiple firms, markets or consumers.

Outside of prudential, systemic and major financial services risks, there are conduct-related and complaints-related risks. Financial services bodies should come together under the Wider Implications Framework, whose members include the Financial Ombudsman Service, the Financial Services Compensation Scheme, The Pensions Regulator and Money and Pensions Service (MaPS), to consider the impact of AI on its cross-regulator work programme.

- b. **Fraud, financial crime and cyber risks:** Continue to strengthen coordination with regulators, government agencies, and the private sector, recognising that the FCA is already a key partner in the UK's whole-system approach to tackling these threats.

As AI increases the speed, scale and sophistication of fraud and cyber-attacks, maintaining this as a strategic priority will become increasingly important. Coordination and collaboration should support earlier intelligence-sharing, identification of common warning signals and faster escalation.

This is particularly important where no single firm, public authority or jurisdiction sees the full picture and risks may propagate more quickly.

Our specific strengthening coordination recommendation in this area relates to engagement with AISI. The FCA should ensure that it has appropriate relationships with AISI, especially in relation to advances in model capabilities that may impact financial services. This may include fraud, financial crime and cyber risks, but could extend to other areas as well.

- c. **Cross-sector consumer, data and competition risks:** There is excellent existing coordination among sectoral regulators who have powers in relation to competition, consumer protection and data protection through the Digital Regulation Cooperation Forum (DRCF) and the UK Regulators Network, as well as bilateral engagement at senior levels and MOUs between the FCA and multiple institutions.

However, we do consider that the FCA and sectoral regulators, such as the ICO, CMA and Ofcom, could strengthen their coordination on the impact of AI on consumer protection, especially in relation to agentic commerce and the consumer interface through the DRCF.

Other areas of focus for coordination across these regulators would include AI-enabled impacts on consumer redress and complaints systems, data governance and privacy, market structure and platform power, and risks affecting vulnerable consumers, access to services, and outcomes that span regulated and unregulated platforms.

The sectoral regulators and DRCF should also consider how to engage other sectoral regulators and consumer bodies where AI-enabled activity goes beyond its core membership or crosses multiple regulatory remits.

- 2. Strengthen international cooperation. The FCA plays a significant role, alongside the PRA and Bank of England, on the international stage and is an active member of multiple international forums.

We recommend that the FCA deepen engagement in international forums, including the International Organisation of Securities Commissions (IOSCO), the Financial Stability Board (FSB), the G7 Cyber Experts Group (CEG) and the Bank for International Settlements (BIS) to build shared understanding of cross-border AI systemic risks and more coordinated supervision.

We also consider that the FCA should strengthen its work internationally through the Global Financial Innovation Network (GFIN) and that this body can help support the FCA in promoting innovation and international competitiveness.

3. Coordinated incident response framework. AI risks are new and evolving. Mechanisms such as the Cross Market Operational Resilience Group (CMORG) are helpful, and the technology-neutral approach of the Authorities Response Framework (ARF) enables the key authorities to respond to sector wide operational disruptions and respond to incidents in other sectors that may indirectly affect the finance sector.

However, the FCA and the authorities should consider developing a coordinated incident response framework for AI-related systemic events. This would develop protocols for responding to AI-related incidents affecting multiple firms or authorities through shared AI dependencies.

As can be seen from recent events, the impact might not be systemic but can be widespread and cross-border, and might require involvement of a wider set of partners than previous incidents. Such a protocol could set escalation criteria, define roles, responsibilities and communication channels, and be tested through scenario exercises involving domestic and international counterparts, where relevant.

Implementation considerations

This can largely be delivered through existing institutional arrangements but depends on sustained leadership and governance.

Clear delineation of responsibilities is essential: the FCA supervises conduct, the PRA prudential resilience, the ICO data protection, and the CMA competition. What changes is that AI-related systemic risks increasingly span these domains and require coordinated action.

The FCA must be able to identify system-wide signals, supported by the supervisory capabilities in priority recommendation 6, which enable earlier detection of cross-firm risks and a more timely coordinated response.

Priority recommendation 3: monitor the transition to autonomous models and adapt regulatory frameworks

Recommendation

The Mills Review recommends that the FCA consider monitoring the transition towards more autonomous AI and adapt regulatory frameworks.

AI will become a core operational capability for firms and agentic finance will emerge for consumers and financial services, moving the human along the autonomy spectrum to Approver or Observer. In this context, decision-making in firms and by consumers becomes more distributed across systems and the pace and speed of AI-enabled financial services innovation is likely to increase.

The FCA will therefore need to continue clarifying how key accountability, governance and consumer protection frameworks apply.

We also recognise, as discussed in the Capability Shift chapter, that many conduct, consumer protection, fraud, financial crime and operational risks associated with AI systems are managed through firms' Model Risk Management (MRM) frameworks. As AI becomes a core operational capability, these frameworks and other governance frameworks operated by smaller firms will need to evolve.

General-purpose and frontier AI models can introduce challenges including opacity, model drift, data bias, hallucinations and emergent behaviours that are not always well addressed through traditional point-in-time validation approaches. Firms will increasingly require more dynamic approaches to model governance, monitoring and assurance, supported by end-to-end controls across the AI lifecycle.

The FCA, alongside the PRA, should support the development of more effective approaches to AI model risk management as it relates to conduct risk (and prudential risks for non-dual regulated firms) and, consistent with Recommendation 4, work with firms, model providers and researchers to understand how advances in AI can improve explainability, assurance, governance and oversight.

Why this matters

Existing accountability and consumer protection frameworks were designed for primarily human-mediated interactions. As AI systems move from Operator and Collaborator towards Approver and Observer modes, evidencing accountability may become more complex.

Firms and individuals remain accountable for outcomes, even where aspects of model behaviour, performance and change sit outside their direct control. For example, a consumer journey may increasingly involve interconnected AI systems operating across a firm and third-party providers to support activities such as lending, financial advice, fraud prevention or customer service. This makes it more difficult to evidence how outcomes were delivered, demonstrate reasonable steps, or identify the source of errors, bias or harm.

Consumer outcomes may become harder to attribute, and expectations such as demonstrating reasonable steps under the Senior Managers Regime become more complex where systems operate with greater autonomy.

As decisions are increasingly delegated to AI systems, it may become harder for firms to deliver and demonstrate consumer understanding, value and informed consent. Dependence on third-party models and infrastructure also introduces challenges for control and assurance.

The issue is therefore how existing frameworks apply as AI moves along the autonomy spectrum and as the debate around the legal personality of AI systems continues. Although Law Commission statements confirm that AI systems do not have legal personality and therefore cannot be sued or prosecuted, engagement with stakeholders points to ongoing debates. By monitoring developments, the FCA will be able to identify where clearer interpretation, and more substantive adaptation may become necessary.

Clarity on how these frameworks apply is also important to support the FCA's secondary objective of promoting the international competitiveness of the UK economy and its growth in the medium to long term.

Providing regulatory certainty on how outcomes-based regimes such as the Consumer Duty apply in AI-mediated contexts can support innovation and inward investment and strengthen the UK's position as an attractive location for the development and deployment of agentic financial services.

Key actions

This reflects increasing complexity at intermediate levels of autonomy and emerging pressure points at higher levels:

1. Monitor the progression along the autonomy spectrum: Track how firms move from Operator and Collaborator uses of AI towards Approver and Observer modes, including where systems act on behalf of firms or consumers and where decision-making becomes more distributed. This should highlight where accountability and consumer protection considerations become more material and where dependencies may require closer supervisory attention.
2. Clarify for firms how existing frameworks apply in AI-mediated contexts: Understand how the Consumer Duty, Senior Managers Regime and governance expectations apply where systems are more adaptive, decisions are distributed, and outcomes evolve. This includes how firms show evidence of outcomes in dynamic journeys, how senior managers demonstrate reasonable steps, and how accountability is exercised across interconnected systems.
3. Strengthen expectations on evidencing outcomes and control: Ensure firms can demonstrate how outcomes are delivered, how controls operate and how accountability is exercised where AI systems are embedded in decision-making. This includes monitoring system behaviour, testing that extends beyond deployment into operational performance, controls that detect when systems move outside expected bounds, and governance processes that respond to issues identified in operation. They should look at evolving model risk management to address model drift, data bias, opacity and other risks associated with increasingly capable AI systems.
4. Assess when further adaptation may be required: Use supervisory evidence to highlight where the interpretation of existing frameworks needs to be more explicit or whether to adapt them as autonomy increases, including the role of regulatory frameworks such as the Critical Third Party or the Designated Activities Regime.
5. Improve the pace and responsiveness of regulatory adaptation: Ensure FCA supervisory, policy and rulemaking processes can respond in a timely and transparent way to the impacts of rapidly evolving AI capabilities and market developments. This should include clarifying regulatory expectations, identifying emerging harms, and adapting relevant regulatory frameworks.

Implementation considerations

This can be delivered within the existing regulatory framework, but effectiveness depends on sustained monitoring and supervisory clarity. Firms will need to maintain accountability as decision-making becomes more distributed, without constraining AI adoption. The FCA should set expectations that are proportionate to the level of autonomy and risk, recognising that firms will progress at different speeds while maintaining consistent consumer protection.

Priority recommendation 4: scale up the FCA's AI Lab to support AI models and system innovation in financial services

Recommendation

The Mills Review recommends that the FCA should consider establishing a structured capability, anchored in the AI Lab, to assess AI models and systems used in financial services. This should include current general-purpose and domain-specific models and systems, while placing strongest emphasis on emerging architectures, more capable models and their potential application to regulated use cases.

Working with firms, model developers, researchers and technical experts, this capability should enable earlier dialogue, shared learning, technical assessment, drawing in practical testing where appropriate. Its forward focus should help the FCA anticipate how new model capabilities and system designs could shape financial services before they become embedded in core firm operations.

This would help the FCA develop an independent understanding of AI models and systems, use its convening role to support responsible growth, and address challenges around explainability, assurance and governance in financial services AI. It would also support the UK's financial services industry as a leader in global AI compliance, assurance and innovation.

Why this matters

AI adoption in financial services is becoming a systemic shift as models move from discrete use cases into core firm operations, including client front-facing operations. Current general-purpose, specialised and domain-specific models are already being deployed, and there is an ongoing need to understand how AI systems perform when governed, integrated and monitored within regulated environments.

The FCA also needs to anticipate emerging architectures and more capable models that may reshape how financial services are designed, delivered and supervised. Earlier engagement with firms, developers and researchers would help ensure that these capabilities are designed and adapted for regulated financial services, where systems must be explainable, auditable and accountable in practice.

The FCA may require independent capability to understand and evaluate model and system behaviour in practice, rather than relying solely on firms' or providers' own assessments.

Building this capability early would support safer adoption of current models, clearer supervisory expectations and readiness for more advanced systems, while contributing to responsible growth, stronger supervision and UK financial services competitiveness.

Key actions

The FCA should consider establishing and operating this capability through a structured, partnership-based approach:

1. **Build a core capability:** Establish an FCA capability, anchored in the AI Lab, to assess AI models and systems in financial services contexts. This should build from current general-purpose, specialised and domain-specific models, while focusing most strongly on emerging architectures, more capable models and their application to regulated use cases.
2. **Engage earlier in the development cycle:** Work with firms, model developers, researchers and technical experts to understand how new capabilities are being developed, adapted and governed for financial services, and to identify opportunities where emerging model capabilities could address current regulatory, governance and assurance challenges. This should help bring regulatory insight to design and deployment choices where this can add value.
3. **Use structured partnerships:** Develop partnerships that provide access to specialist expertise, independent challenge and realistic use cases, while helping the FCA build its own understanding of model and system behaviour.
4. **Set initial priorities:** Focus early work on emerging model capabilities and system designs, and how these may need to be adapted for regulated financial services. This should include understanding how advances in model architecture may improve explainability, auditability, assurance, monitoring and governance in regulated environments. This should complement existing testing work by developing a more upstream view of model behaviour, assurance needs and design choices as they evolve.
5. **Publish practical outputs:** Share insights and lessons to inform supervision, firm engagement and dialogue with model providers, supporting responsible adoption today and readiness for more advanced systems.
6. **Review and scale:** Assess whether the capability is improving FCA understanding, supporting safer adoption and strengthening market practice, then adjust scope, partnerships and resourcing as AI use develops.

Implementation considerations

This capability should remain targeted and pragmatic, focused on producing usable outputs rather than building large-scale infrastructure. It should inform firms and supervision without becoming a certification function or duplicating existing testing initiatives. Whether it is effective will depend on maintaining independent analysis while ensuring access to real-world use cases and technical expertise.

Priority recommendation 5: enable the foundations for agentic finance

Recommendation

The Mills Review recommends that the FCA should consider leading the development of a trusted framework for AI agent participation in financial services, clarifying how agents can be authorised, identified and held accountable. This should help establish clear expectations for consent mandates, identity, control and liability, creating the conditions for safe adoption of more autonomous AI-enabled services.

Why this matters

Financial services are beginning to move along the autonomy spectrum from assistive and co-worker roles towards Approver and Observer modes. This presents a significant opportunity to transform how consumers access and use financial services, particularly through AI agents that can support decision-making, execution and ongoing financial management.

If adopted safely and at scale, agentic finance has the potential to expand access, help address the advice and protection gaps, and enable more efficient and personalised business models.

This could reduce friction in financial services journeys including onboarding, lending, mortgages, protection, portfolio management for consumers, small businesses and corporates. This can lead to both the emergence of new products and innovation but also a general boost to productivity and economic growth if agents can be trusted to operate safely in financial services.

The opportunity is vast, but progress is currently significantly constrained by fragmentation across data access, identity verification for agents, consumers, businesses and payments infrastructure. This limits agents' ability to operate reliably and safely end-to-end.

These components need to function as an integrated system. An agent requires access to data, a trusted identity, and to be able to execute actions to complete tasks. Where these elements are not aligned, systems cannot safely progress beyond assistive roles into coordinated or delegated activity.

Likewise, progress on digital identity for individuals and businesses is critical to underpin agentic finance and an agentic economy in the UK. While this is outside of the FCA's remit, it is a critical component or enabler for the benefits of an agent economy to be fully realised.

This creates a near-term policy consideration. Without clear participation frameworks at earlier stages of adoption, accountability and control risks may begin to emerge before more autonomous use becomes established. Decisions taken now on how agents are authorised, identified and held accountable will shape how systems progress.

Establishing clear frameworks for agentic finance also supports the FCA's secondary objective of promoting the international competitiveness of the UK economy and its growth in the medium to long term.

By providing regulatory certainty at an early stage, the UK can position itself as a leading jurisdiction for the development and deployment of agentic financial services, enabling firms to innovate, scale new business models and attract inward investment ahead of other markets. A trusted framework for AI agents can support safe adoption while allowing the UK to capture the benefits of this transition.

Key actions

The FCA should consider:

1. **Defining requirements for AI agents acting in financial services:** Establish minimum conditions for those acting on behalf of consumers, including valid instructions, structured and verifiable mandates (scope, limits and revocation), and requirements for identity linkage, reflecting increasing autonomy.
2. **Clarifying accountability and liability frameworks:** Set expectations for how existing frameworks apply, including the Consumer Duty and senior managers regime, ensuring clear allocation of responsibility as decision-making becomes more distributed.
3. **Developing standards for evidencing control and auditability:** Define how firms should evidence agent behaviour, control operation and outcomes, including expectations for monitoring, testing in live operation and intervention where systems operate outside defined parameters.
4. **Enabling interoperability across data, identity and payments:** Support alignment with ongoing open finance and payments developments so that agents can operate across systems in a way that enables progression while maintaining security and control.
5. **Engaging domestically and internationally on agent frameworks:** Work with other regulators and international partners on interoperable approaches, recognising that progression toward Observer mode, that is, more autonomous AI (Level 5) will require consistency across sectors and jurisdictions.
6. **Engaging with government on the supportive enablers for trusted agent frameworks** including digital identity for individuals and how this interacts with the full capability of an agentic finance system.

Implementation considerations

This should be developed as a targeted, principles-based framework, built on existing regulatory foundations rather than creating a new regime. The objective is to enable adoption while ensuring that accountability remains clear and enforceable.

The framework can be developed through a combination of approaches, depending on the desired pace of adoption, interoperability and market penetration.

There are four options:

1. **Develop a trusted agent protocol.** The first is to build a trusted agent protocol that firms would need to follow as they build AI agents, or it could be an agentic finance protocol.
2. **Develop trusted agent standards in Open Finance (Recommended).** Use the FCA's work on Open Finance as a means to build standardised approaches within the context of data sharing and adjacent actions for AI agents acting with Open Finance data.

A separate industry committee could be set up to work on this with regulatory input and build standards that have the benefit of meaningful penetration within the financial systems ecosystem, cross markets. The principal disbenefit is timing. Open Finance may take longer to mature than the pace of AI adoption, meaning agentic use cases beyond payments could emerge before common market standards are established across areas such as savings, investments, insurance and lending.

We recommend this option given the Open Finance roadmap which has been set out for the FCA and is underway, and the need to move beyond agentic payments to consider a wider set of agentic products and services. We recognise that options 1 and 2 could work sequentially or together.

3. **Work on a trusted agent protocol internationally.** The third is to work with international standard-setters on principles. This approach could, of course, be combined with domestic work. The FCA often joins a range of international committees to work on general principles of wider application while working on more detailed components within the context of its domestic obligations. Working in this way, both within the domestic setting and internationally, would go towards addressing concerns raised by firms in engagement responses around interoperable standards, both within the UK and cross-jurisdictionally.
4. **Work on a trusted agent protocol cross-sectoral regulators or cross-economy.** A bolder approach would be to establish work across regulators, either through the Digital Regulation Cooperation Forum (DRCF) or another combination of sectoral regulators, to build a trusted agent protocol cross-economy. This could link with the government's smart data work and with the work of the Smart Data Council. The benefits of a protocol for the UK sectoral regulator or the UK economy could be vast, provided it were interoperable, and sufficiently broad to encompass multiple consumer and business journeys. The Review recognises that this might go beyond the powers or the capacity of the FCA.

Overall, the Review considers Options 1 or 2 provide the most viable approaches for developing trusted agent frameworks in financial services, either sequentially or in parallel.

Priority recommendation 6: build and adopt an AI-enabled agentic supervisory model

Recommendation

The Mills Review recommends that the FCA consider developing an AI-enabled agentic supervisory model that enhances both firm-by-firm regulatory efficiency and system-wide oversight.

This includes improving effectiveness across the full regulatory lifecycle, from authorisation, supervision and enforcement, while building capability to detect new kinds of risk.

This builds on, rather than replaces, human-led regulatory judgement and accountability and existing firm-based supervision models. Supervision should support more effective firm-level engagement and more continuous, system-wide oversight so that harm can be detected and acted upon more quickly.

It will give the FCA a stronger ability to speed up authorisation and support firms to grow and innovate at the gateway and in its Scale-up Unit. This will help deliver its secondary objective on growth; detect firm-level compliance issues and risks of harm; and, crucially, monitor system-wide risks across firms through shared models, data, intelligence and infrastructure.

Why this matters

AI-enabled supervision serves two related purposes.

First, it can make regulation more effective across the full regulatory lifecycle, with human supervisors focusing on areas where judgment, discretion and engagement matter most, while maintaining oversight and ensuring the FCA's own use of AI meets the standards it expects of firms.

Second, the agentic supervisory model addresses system-wide risks that firm-by-firm oversight cannot detect. Financial services firms increasingly rely on shared models, cloud platforms and data sources. Changes in these systems can affect multiple firms simultaneously, with impacts that may not be visible within any one firm, until they evolve into industry-wide patterns of harm.

Current supervisory approaches have largely been designed for risks within individual firms and identified through periodic review. That remains important but may no longer be sufficient on its own. Some harms only become discernible when viewed across the market, such as emerging fraud, financial crime or conduct risks.

Where a widely used model changes behaviour through updates or data drift, firms may observe only marginal signals, insufficient to prompt action. The underlying pattern sits outside any single firm's view, and by the time it is identified, harm may already have occurred at scale.

An AI-enabled regulatory model for supervision could allow the FCA to observe outcomes across firms in near real time, detect anomalies quicker, connect patterns that individual firms cannot detect, and intervene earlier. This supports better consumer outcomes, maintains market integrity and strengthens confidence in AI-enabled financial services.

Key actions

The FCA should consider using a phased approach to enable an agentic supervisory model:

1. Enhance firm-based supervisory efficiency: Develop and deploy AI-enabled tools for authorisation, supervision, enforcement and administrative processes. They should also support supervisors in identifying risks, prioritising activity and analysing information, and automate routine, lower-risk processes.

On our autonomy spectrum, the FCA should develop AI to the Consultant and Approver levels, not Observer. This would mean that human supervisors remain responsible for key decisions including exercising judgement, overseeing AI usage, and engagement with firms (outside agent-to-agent communication).

2. Prepare for agentic supervision and agent to agent workflows:

For a financial services regulator, agentic AI creates the opportunity to move from a supervisory model that is largely episodic, document-based and human-mediated, towards one that is more continuous, risk-based and intelligence-led.

Today, supervision often depends on firms interpreting obligations, assembling evidence, submitting returns or reports, and responding to follow-up questions, while supervisors review that material, compare it with expectations and decide whether further engagement is needed.

Agentic AI provides an opportunity to reform this operating model by enabling supervisory agents to support more of the end-to-end process: triaging firm submissions, testing evidence against supervisory expectations, generating information requests, maintaining issue logs and supporting earlier intervention.

This could make existing processes faster and help develop a more responsive supervisory model that can keep pace with an increasingly AI-enabled financial services industry.

The autonomy spectrum provides a way to illustrate examples of the different types of tasks that could take place as the agentic supervisory model matures:

- Operator and Collaborator level: agents support supervisors as assistants in searching, summarising, drafting and checking information while humans remain fully in control.

- Consultant level: agents begin to coordinate supervisory workflows in monitoring risk indicators, linking evidence within and across firms, proposing supplementary enquiries, and consulting human supervisors at judgement points.
- Approver level: agents may prepare supervisory actions, information requests or thematic findings for supervisor approval, while continuously monitoring and escalating outliers or exceptions to supervisors within clear guardrails.

This is not a comprehensive vision of the model but an example of how an agentic supervisory model could provide speed, scale and workflow orchestration while human supervisors remain in the loop and retain responsibility for judgement and decision-making.

3. Agentic models for system-wide oversight: Build capability and tools for system-wide oversight, including strengthening technical capability to understand how AI systems operate, and how risks emerge and propagate across firms, systems, shared models and infrastructure.

Develop AI-enabled tools that monitor outcomes across firms in near real-time and can detect patterns and risks not visible at firm level. This could include market surveillance tools that analyse transaction and order book data across markets to identify potential market abuse more quickly; tools that detect system-wide risks arising from common exposure to AI models, data providers or infrastructure; aggregated monitoring of complaints, pricing and Consumer Duty outcomes across firms to identify harms that may not be visible within individual firms. It also includes horizon-scanning capability that analyses regulatory data, market intelligence, promotions and other information sources to identify emerging risks, misleading promotions, fraud typologies and other cross-market patterns earlier.

These capabilities require firms to provide timely, structured data on outcomes, and would combine data, AI models and agentic workflows to identify risks and surface actionable insights for earlier supervisory intervention.

Supervisory judgement remains central: AI enables earlier identification, with decisions on intervention and enforcement taken by human supervisors.

4. Adapt the operating model for faster action: Operate supervision, policy, intelligence and enforcement in a more coordinated way, so that insight translates rapidly into action. This includes moving from periodic reporting towards more continuous, event-driven data and intelligence flows that improve timeliness while reducing the burden on firms from ad hoc data requests. This should support both firm-level oversight and system-wide pattern detection.
5. Develop FCA workforce capability, governance and operating models for an AI-enabled regulator. This should include strengthening AI literacy, supervisory and technical capability; updating accountability and governance frameworks; and adapting roles, responsibilities and ways of working to support the automation and delegation of regulatory workflows and tasks, while maintaining human judgement, oversight and accountability for regulatory decisions.

AI adoption will require significant workforce upskilling, evolving operating models and management of workforce transition to realise benefits safely and effectively.⁸⁶

Implementation considerations

Delivering this model depends less on technology alone and more on how supervision adapts, and how insight is generated, triaged and acted upon.

It requires investing in agent-to-agent protocols tailored to a regulatory environment, agent identity and permissions, audit and escalation protocols. A key investment would be the data and semantics layer of the stack and we would recommend a deep dive into this area to underpin feasibility for widespread supervisory agent to firm agent communication across the agentic supervisory model.

The challenge is whether or not machine-readable data is a key element for widespread agent-to-agent communication capabilities in the model, or whether, for risk management and responsible adoption, the plan might need to work through agent-to-agent communication on a case-by-case basis.

Moving towards more continuous data flows improves timeliness and reduces duplication, but requires firms to provide more structured, high-quality data and may challenge existing reporting models. Implementation should therefore be phased and proportionate, focusing on the biggest risks.

Implementation is also likely to vary across the FCA's regulated population. Larger firms may be able to adopt agent-to-agent workflows and structured data approaches sooner, while smaller firms may require simpler interfaces, longer transition periods or alternative engagement models.

As supervisory capability becomes more continuous and data-driven, the FCA will also need to consider how supervisory tolerances, materiality thresholds and risk appetite are calibrated so that increased monitoring capability does not inadvertently discourage innovation, competition or proportionate risk-taking.

Firm-level oversight remains important but must be complemented by system-wide visibility.

Priority recommendation 7: develop a trusted public-interest AI-enabled financial capability service

Recommendation

The Mills Review recommends that the FCA consider taking a more proactive, 'hands on' approach to financial capability and access to future vital AI-enabled services.

The FCA should consider convening the development of a free, inclusively designed, trusted AI-enabled financial capability and support system that has a range of digital tools that provides consumers with access to reliable financial information, guidance and support from trusted sources.

⁸⁶ [FPS: the future of AI and the workforce](#), City of London, 2025

As AI increasingly becomes a primary way through which consumers access information, compare options and make financial decisions, access to high-quality support should not depend solely on consumers' ability to pay for the most capable AI models or services.

The FCA should consider working with government, MaPS, consumer bodies and industry to explore the feasibility of a public-interest, or sovereign-style, AI-enabled financial capability service that is free at the point of use and designed to support better consumer outcomes.

Delivery should be through broad-based partnership, convening partners across the regulatory and consumer advice landscape to collaborate and share responsibility for design and delivery.

Why this matters

Financial capability and access to products and services are essential for good customer outcomes in financial services. By 2030, AI will increasingly shape those outcomes.

AI brings a significant opportunity to tackle the advice, debt management and protection gaps in UK financial services by expanding access to financial information and support at lower cost and greater scale. If deployed responsibly, AI-enabled support could help consumers make more informed decisions regarding savings, pensions, mortgages, debt, insurance and protection products.

However, there is a risk that these benefits are unevenly distributed. Access to the most capable AI systems may increasingly depend on cost, digital capability, commercial arrangements or platform choice. Consumers who would benefit most from financial support may therefore face the greatest barriers to accessing it.

Should frontier AI capabilities become concentrated within paid services or services with embedded commercial interests, unequal access and conflicts of interest may emerge.

To help markets work well for consumers, the FCA should therefore work with industry, government and other regulators to support the provision of timely, accurate and impartial financial information and support through AI-enabled tools, helping ensure consumers can access free information that supports their financial health regardless of income or ability to pay for commercial AI services.

As highlighted in the consumer research and consumer chapter of this Review, consumer trust is critical. Consumers are already using general-purpose AI tools to support financial decisions. Without access to trusted, high-quality sources of information and support, consumers may either disengage or rely on unverified services. This may lead to poorer outcomes and reduced confidence in AI-enabled financial services.

In securing an appropriate degree of protection for consumers, the FCA should have regard under section 1C of FSMA to, among other matters, the differing levels of experience and expertise that consumers may possess, and the need consumers may have for timely information and advice that is accurate and fit for purpose.⁸⁷ As AI increasingly mediates access to financial information and support, these considerations become more important.

We recognise that this must be balanced against the principle that consumers remain responsible for their own financial decisions. The objective is not to remove accountability from consumers, but to ensure that free, fair and impartial information is available to support them in making informed and accountable choices, rather than having decisions shaped by inaccurate, conflicted or inaccessible sources of information.

A regulatory approach that enables broad access to a free, inclusively designed and trustworthy financial capability tool would help mitigate these risks and support more equitable realisation of AI benefits across the population.

Key actions

The FCA should consider taking a structured, partnership-based approach.

1. **Develop the concept for a trusted AI-enabled financial capability service:** Work with MaPS, HM Treasury, consumer organisations, regulated firms and technical experts to define the objectives, scope and core design principles of a trusted AI-enabled financial capability service that provides consumers with free access to reliable financial information, guidance and support. This should include consideration of a sovereign-style or public-interest model built on trusted financial information sources.
2. **Convene relevant stakeholders:** Identify and convene organisations required to develop, govern and deliver such a tool, including HM Treasury, the Money and Pensions Service, Financial Ombudsman Service, Financial Services Compensation Scheme, regulated firms, trade associations, consumer organisations and the debt advice sector.
3. **Assess feasibility and delivery models:** Work with stakeholders to assess feasibility, appetite and potential barriers to delivery, recognising that incentives may not be fully aligned and that the public benefit will need to be clearly articulated. This should include consideration of whether the capability should be delivered through a sovereign-style model, a public-interest model, integration with existing services, or through a partnership-based approach using trusted external providers.
4. **Define responsibilities and resource requirements:** Establish roles, responsibilities and funding models across participants to support delivery and long-term sustainability.

⁸⁷ Financial Services Market Act, 2000, <https://www.legislation.gov.uk/ukpga/2000/8/section/1C>

5. Develop delivery options: Explore practical delivery models, including integrating existing services into a unified interface, commissioning a third-party solution, or industry-led delivery supported by public-interest governance arrangements.
6. Enable implementation through partnership: Support the development and rollout of the service, ensuring it provides clear pathways to key services such as debt advice, pensions guidance, savings support, consumer redress mechanisms, money management and fraud and scam awareness.

Implementation considerations

The FCA could act on its own, depending on available resource and on interpretation of its existing powers. However, a broad-based partnership approach is likely to deliver stronger results.

It must coordinate with government, including HM Treasury, and carefully manage engagement with industry to address competition and governance issues. Existing public-interest bodies such as MaPS may be well placed to play a central operational role, with the FCA acting as convenor and strategic sponsor.

The purpose of the service is to provide consumers with access to trusted, impartial and high-quality information and guidance in an increasingly AI-mediated financial system.

The scope of the service should be carefully defined and aligned with known consumer priorities. This could include, for example, debt advice pathways, access to redress, pensions and savings guidance, and information on spotting and avoiding fraud and scams.

Over time, this capability could become part of the wider foundations for an AI-enabled financial services ecosystem, ensuring that the benefits of increasingly capable AI systems remain accessible to all consumers rather than only those able to access commercial frontier AI services.

Annex

- I Glossary
- II Engagement paper responses
- III Consumer survey findings
- IV Consumer focus group findings
- V Economic growth impacts
- VI International approaches to AI
- VII Infrastructure for agentic finance

Annex I - Glossary

Given the emerging nature of AI in financial services, we have defined key terms for the purposes of this report. These terms may be used differently elsewhere, and the definitions set out in the Mills Review should not be taken to represent a fixed or definitive FCA position.

Advice-like support: Highly personalised support provided by AI frontier models that falls outside the FCA's regulatory perimeter. Such support would be treated as regulated financial advice if it fell within the scope of the perimeter.

Agentic AI: AI systems composed of one or more co-ordinated AI agents that can autonomously break down tasks, collaborate and pursue complex objectives over extended periods in dynamic environments, with minimal human supervision.

Agentic finance: Finance in which at least one party to the allocation, intermediation or use of capital is an autonomous AI agent capable of independent economic action.

AI agent: A system that can perceive and act upon its environment with a degree of autonomy, using tools as needed to achieve specific goals and adapt to changing inputs and contexts.

AI-enabled system: A system that uses AI to generate outputs, inform decisions or take actions in support of defined functions, tasks or objectives.

AI governance: The structures, policies, processes and controls through which an organisation directs, monitors and accounts for the design, deployment and use of AI-enabled systems.

Artificial General Intelligence (AGI): AI that can perform a wide range of economically valuable cognitive tasks at or above human capability, rather than being limited to narrow or specific functions.

Artificial intelligence (AI): Systems that infer from data to generate outputs, such as predictions, recommendations or decisions, and to perform tasks associated with human intelligence, including pattern recognition and language processing.

Assistive AI: AI-enabled systems that support human judgement or action by generating information, suggestions or outputs while leaving decisions and responsibility with the human user.

Autonomous AI: AI-enabled systems that can plan and carry out actions within defined objectives or constraints with minimal or no real-time human input.

Bias: The risk that AI systems may produce systematically unfair or discriminatory outcomes due to biased training data, model design choices or deployment context.

Cognitive erosion: A reduction over time in a person's knowledge, judgement, confidence or decision-making capability associated with sustained reliance on AI-enabled systems.

Continued influence effect: The tendency for misinformation to continue shaping a person's thinking or decisions even after it has been corrected.

Cryptography: Techniques used to secure data and communications by keeping information private, preventing unauthorised changes, and ensuring actions and transactions can be trusted and verified.

Deceptive design: Interface or service design that steers, pressures or misleads people into choices that may not reflect their interests, preferences or informed consent.

Delegation: The extent to which decision-making or action is assigned from a person or organisation to an AI-enabled system within defined objectives, permissions or constraints.

Delegation levels: A five-level framework describing increasing delegation to AI-enabled systems (see [cross-reference to relevant section]).

Deployer: An organisation, often a regulated firm, that implements and uses an AI system within its operations or consumer services, including configuring, integrating, monitoring and governing it.

Disintermediation: The reduction, removal or transformation of intermediary roles in a market or process, enabling more direct interaction between consumers, firms, platforms or systems.

Explainability: The extent to which the functioning of an AI system can be technically understood, including how inputs are used and how outputs are generated, as distinct from whether a person can make sense of it (see Intelligibility).

Foundation model: A general-purpose AI model trained on broad data that can be adapted or deployed across a wide range of tasks and use cases.

Frontier model: A model at the cutting edge of current AI capability, typically among the most capable available and often requiring the greatest resources to train.

General-purpose LLMs: Large language models with broad capabilities that can be used or adapted across many different tasks, services and downstream applications.

Generative AI (GenAI): A subset of AI systems capable of generating text, images, code or other content based on learning from training data.

Guardrails: Technical or procedural controls designed to limit unsafe, inappropriate, unauthorised or unreliable AI outputs or actions.

Hallucination: When an AI model produces an answer that sounds plausible but is false, unsupported or invented.

Herding: System-wide effects that occur when many firms or consumers use similar models, data, tools or optimisation strategies, leading to aligned behaviours and outcomes that may amplify risk or reduce resilience.

Human-in-the-loop: A design approach in which a human must review or approve AI outputs before they are acted upon, making the human the final decision-maker in the process.

Human oversight: The requirement that AI systems be subject to human monitoring and supervision, enabling individuals to understand system behaviour, detect issues and intervene to reduce risks to people or markets.

Inference: The stage at which a trained AI model is used to generate an output such as a prediction, recommendation, decision or action.

Intelligibility: The extent to which a person can make sense of what an AI system has done, what it means for them and what they can do about it, as distinct from the technical explainability of the system itself (see Explainability).

Meaningful human control: The ability for humans to exercise real understanding, authority and intervention over AI systems, such that they can monitor, challenge and override automated outputs and remain accountable for outcomes.

Multimodal model: An AI model that can work with more than one type of input, for example text, images, audio, video or data.

Neuro-symbolic: A hybrid approach to AI that combines neural networks that learn from data with symbolic reasoning that uses rules, logic and knowledge.

Open-weight model: An AI model whose internal parameters are made available for others to host, adapt or inspect more directly than a closed commercial model.

Open Finance: The extension of secure data sharing and portability principles beyond Open Banking to a wider set of financial products, such as savings, investments, pensions and insurance, enabling more joined-up services and consumer-led data use.

Quantum: Technologies that exploit quantum effects, the behaviour seen at very small scales that differs from everyday physics, for computation and communication, with potential to enhance modelling and optimisation while creating risks to how data and transactions are secured and verified over time.

Annex II - Engagement paper responses

This annex summarises the key findings from 140 stakeholder responses to the engagement paper. The responses broadly support the report's analysis but highlight important tensions and uncertainties across each of the six thematic areas. Overall, stakeholders recognise the transformative potential of AI in retail financial services, while emphasising that outcomes will depend on governance, market dynamics and consumer capability rather than technology alone.

1. Technological change and capability

Respondents broadly agree that recent advances in AI, particularly foundation models and emerging agentic systems, represent a step change in capability. Many firms are already adopting AI tools across a wide range of functions, with the most common applications in customer service, fraud detection, data analysis and internal processes. While adoption is widespread, it remains uneven, and higher-risk, customer-facing uses are still relatively limited. A central theme in responses is the shift from AI as a tool that supports decision-making to systems that can act within workflows and potentially on behalf of users. Stakeholders see agentic AI as a critical development, but also one that introduces significant uncertainty. But respondents highlight that these systems are often probabilistic, difficult to fully understand and capable of producing outputs that are plausible but incorrect. As a result, there is a widely held view that firms will continue to adopt AI cautiously in higher-stakes contexts until confidence in reliability, control and assurance improves. At the same time, there is an expectation in parts of the industry - particularly among around 41% of technology providers and more digitally advanced firms - that capability could improve quickly and lead to more rapid deployment in selected use cases. This creates uncertainty not only about the eventual end-state, but also about the pace at which different parts of the market could move. Respondents also saw fraud and financial crime as areas where AI could change the threat landscape and not just firm capabilities, with risks becoming fast-moving, more complex and harder for firms and consumers to detect.

2. Firm adoption and business model change

Stakeholders anticipate significant changes to how financial services firms operate, although views differ on the pace and extent of transformation. There is strong agreement that AI will reduce costs across operational processes, including customer servicing, underwriting, compliance and software development. Respondents also expect firms to increasingly rely on third-party model providers and vendors, with AI capabilities becoming more accessible through off-the-shelf tools. At the same time, firms highlight that adoption involves substantial investment in integration, governance and assurance. Many respondents note that deploying AI safely and effectively requires significant organisational change, including new skills, revised governance frameworks and enhanced oversight. This suggests that while AI may reduce some costs, it will also introduce new ones, particularly in relation to risk management and system design. However, there are some key risks that were highlighted by specific groups for example Financial Services firms expressed concerns around AI and Cloud Concentration on one or another vendor from a small subset which means the optionality of choices is greatly reduced.

Regulatory and accountability gaps especially with fast pace of AI as well as unclear or lack of traceability within AI systems decisions were also considered as key risks.

Responses also point to changing business models. AI could reduce reliance on traditional intermediaries and enable more direct distribution, but there is uncertainty about how far this will go. Autonomy within AI decisions especially in advanced market models of agentic decisions was flagged as a key risk by some technology as well as Financial Service firm responders. Several stakeholders emphasise that trust, liability and regulatory requirements will continue to shape how services are delivered, and that many existing structures may persist, albeit in modified form. Overall, respondents expect meaningful shifts in how firms produce and distribute financial services, although views differ on whether these changes will be gradual or more rapid in particular markets or use cases.

3. Consumer journeys and outcomes

The responses highlight both the benefits and the risks of AI-enabled consumer journeys. On the positive side, stakeholders see significant opportunities to improve consumer outcomes. AI could help consumers better understand products, compare options, manage their finances and access guidance. There is particular potential to address longstanding issues such as the advice gap and to improve outcomes for consumers with lower financial capability.

But these risks are not hypothetical. Consumer advisers and advocacy groups report multiple cases of harm already occurring, including individuals making poor financial decisions based on inaccurate or incomplete AI-generated information. A consistent concern is that consumers may over-trust AI outputs, particularly where they are presented confidently, and may struggle to identify mistakes or inaccurate information. This creates clear risks of detriment, particularly where consumers rely on AI without independent verification.

Delegation is a critical issue. As AI systems move from providing information to acting, it complicates the relationship between consumers and financial services. Respondents highlight concerns about consumers losing agency, understanding, and not being able to tell when things have gone wrong. These risks are compounded as consumers become increasingly dependent on AI over time.

Inclusion and exclusion emerge as central themes. While AI could broaden access to financial services, it could also widen existing inequalities. Stakeholders warn of a potential two-tier system, in which consumers with access to high-quality AI tools get better outcomes than those who cannot. Digital capability, trust and affordability are determining factors.

Several respondents also raise concerns about the potential erosion of traditional risk pooling in markets such as insurance. More granular data and personalised pricing enabled by AI could weaken cross-subsidisation between consumer groups, potentially improving accuracy and efficiency but reducing affordability or access for higher-risk consumers.

This may create new distributional challenges and raise wider questions about fairness and inclusion in essential financial services.

Respondents also raised the steep increase in AI-enabled deception and the evolution and escalation of this as we move toward 2030 including cross platform scam orchestration, synthetic data misuse and model over reliance when new typologies emerge.

4. Competition and market dynamics

Responses indicate that AI is likely to reshape competition in retail financial services, although there is significant divergence in views on how this will play out. There is a widely shared view that control of the customer interface will become increasingly important. As consumers use AI tools, platforms and agents to search for, compare and select financial products, the ability to influence these interfaces may become a key source of market power.

However, respondents disagree on whether AI will increase competition or reinforce concentration. Some stakeholders emphasise that AI could lower barriers to entry by reducing the cost of building and scaling financial services businesses, enabling new entrants to compete more effectively. Others argue that scale advantages, particularly in data, capital, infrastructure and regulatory capacity, will continue to favour large incumbent firms. This divergence is particularly evident in views on whether efficiency gains will be passed through to consumers or retained within firms. This has been further expressed as one of the key concerns by Financial Services firms so that consumer vulnerability especially lack of awareness or proper technical training for consumers means they may not be capable of sufficiently and safely interacting with modern AI systems.

There is also growing concern about upstream concentration. Many respondents highlight the reliance on a small number of model providers, cloud platforms and infrastructure providers, which may shape pricing, access and innovation across the sector. Even where downstream competition remains active, upstream dependencies may influence market outcomes in important ways.

Intermediation is expected to evolve rather than disappear. While some traditional intermediaries may face pressure, respondents anticipate the emergence of new forms of intermediation, particularly through AI platforms and agents. This could shift how value is captured within the sector and change how consumers access financial services.

5. System-wide and macroeconomic effects

Stakeholders generally had fewer concrete views on the broader economic implications of AI-enabled financial services, but several themes emerge. Respondents recognise that changes in consumer behaviour, such as more active switching and more efficient allocation of savings, could have wider effects on capital allocation and financial stability.

Reduced inertia in savings and investment decisions may improve efficiency but could also affect firm funding models and the stability of certain product markets.

There are also concerns about herding behaviour, where similar AI services and models lead to more aligned consumer decisions. This could increase the speed and scale of market movements and amplify volatility in certain contexts. While some respondents note that similar dynamics have occurred with previous innovations, the potential for AI to operate at scale and at speed is seen as a distinct feature of this cycle.

More broadly, stakeholders acknowledge that the impact of AI on the wider economy, including on employment, incomes and inequality, will feed back into retail financial services. Changes in household financial resilience, income distribution and economic activity will influence demand for financial products and services over time. There is around 75% certainty amongst respondents that market fragmentation would occur and further niches would emerge. This could be positive for the newer market entrants to take advantage of but may also require regulation to adapt to changing technology faster.

6. Regulatory framework and oversight

Respondents broadly support the continued relevance of existing regulatory frameworks, particularly outcomes-based approaches such as the Consumer Duty. However, there is a relatively strong view that these frameworks will become more difficult to apply in AI-mediated contexts. In particular, stakeholders highlight challenges in evidencing consumer understanding, fair value and appropriate outcomes where decisions are dynamic, personalised or partially automated.

A key issue identified is the regulatory perimeter. AI systems and platforms are increasingly able to provide information, guidance and decision-making support that resemble regulated activities, but may fall outside the existing perimeter. This creates concerns about uneven regulation and potential gaps in consumer protection, particularly where consumers interact with general-purpose AI applications.

Another common theme is the tension between accountability and control. While firms expect to remain responsible for outcomes, many report that they have limited visibility into or control over the behaviour of upstream models and systems. This creates challenges for governance, assurance and regulatory oversight, particularly where important inputs sit outside the firm. Within the same vein, around 71% of firms had strong negative sentiments around the absence of a data feedback loop when asked about ways of improving AI competition. Respondents also raised the dual challenge of the need to move quickly to detect and disrupt emerging harms whilst also maintaining appropriate governance and accountability.

Finally, respondents emphasise the need for coordination across regulators and jurisdictions. AI-related risks cut across conduct, fraud, competition, data and resilience, and may not be effectively addressed within a single regulatory framework.

Stakeholders highlight the importance of clear guidance, ongoing engagement and adaptive approaches to ensure that regulation remains effective as both the technology and market structures evolve.

Overall synthesis

AI has the potential to deliver significant benefits in retail financial services, including improved efficiency, expanded access and better consumer outcomes.

However, these benefits are not guaranteed. Stakeholders highlight that risks are already materialising in practice, that market outcomes may diverge depending on competitive dynamics, and that the pace of change may be uneven across firms and use cases.

Annex III - Consumer survey findings

As part of the Mills Review, Yonder Consulting conducted a nationally representative survey to understand consumers' current awareness, usage and perceptions of AI when dealing with their finances. The outcomes are unique as they provide fresh insight on attitudes to future applications and how far people say they are prepared to delegate their finances to AI systems as the technology is embedded.

Full details of the quantitative survey can be found here - [Mills Review Yonder Consulting Survey Results](#).

It was an online survey among 5,026 UK retail financial services consumers, defined as individuals holding a day-to-day bank account, such as a current or savings account. This audience represents the vast majority of the UK adult population, with 98% holding such accounts. Fieldwork was conducted via the YonderLive online panel between 21 and 29 April 2026. Quotas were applied to ensure the sample was representative of UK retail financial services consumers across key demographic variables, including age, gender, ethnicity, region, housing tenure, and internet ability.

AI is widely visible but not yet widely trusted with money

AI is now part of everyday life for many consumers, and for some this has translated into use in personal finance. Of the 92% of consumers who report undertaking at least one personal finance activity (e.g. budgeting, saving, paying bills or investing), 16% already report using AI to assist them. This rises to 23% among consumers who have used AI in their work, studies or personal lives. The profile of those who use AI in financial services tends to mirror that of AI users more generally: they tend to be younger, employed, have higher incomes, be from ethnic minority backgrounds, and be more trusting of AI than non-users.

Currently, AI is being used mainly as an assistive tool. It is primarily used to summarise, explain, simplify and compare information. However, some consumers are already taking this a step further. Among those who report using AI for personal finance or financial products, one in five report uploading personal data or documents (e.g. bank statements) to AI services. In addition, 13% of all consumers say they would be willing to give AI real-time access to their banking and financial data. This willingness increases alongside AI usage: 31% among those who use AI daily for any purpose, and 36% among those who already use AI for personal finance or financial products.

Consumers were shown a series of conceptual AI-enabled financial service propositions designed to reflect plausible near-future applications of AI in finance. This found 20% of consumers were likely to use AI that acts autonomously within pre-set goals. Among existing AI users 28% thought they would definitely or probably use this type of AI for financial services.

Currently there is a sizeable group of consumers who are resistant to AI-enabled financial services. Around a quarter to a third of consumers express consistently negative views across a range of propositions. This reflects low perceived benefit, limited comfort with AI, and concerns about accountability and consumer protection.

Consumers see value in using AI where financial decisions are more complicated

AI use varies considerably depending on the type of financial behaviour involved.

Among consumers who review investments, 20% report using AI services, compared with 17% of those managing debt and 16% of those involved in tax planning. By contrast, AI use is lower among those dealing with bills and subscriptions (7%) or managing savings (12%).

Consumers are not, for the most part, using AI to delegate financial decisions. Instead, they use it to navigate complex or unfamiliar decisions while retaining responsibility for the final choice. Activities such as debt advice, investing, pensions and tax planning are areas where consumers have traditionally sought guidance, making AI's on-demand support particularly attractive.

Trust in AI is selective and rarely unconditional

Consumers appear willing to trust AI in areas where its strengths feel intuitive. Among AI users, 75% trust AI to explain things clearly and 71% trust it to provide useful guidance. However, trust falls sharply when it comes to avoiding misleading information (23%) and handling personal information responsibly (21%).

Consumers frequently raise concerns around the accuracy of outputs, reliability of recommendations and the handling of personal data. Just over a quarter (26%) agree that tools such as ChatGPT provide reliable financial information or advice, compared with 36% who disagree. Confidence is considerably higher among those already using AI in financial services, rising to 57%. Even among users, trust is conditional and many emphasise the need for human oversight.

A small group appear more willing to rely on AI outputs uncritically, particularly where they trust it to avoid misleading information. These consumers are more open to autonomous applications but may be less attuned to risks and safeguards.

Concerns reflect wider issues of accountability and protection

Consumers express consistently high levels of concern about AI in financial services. Around three fifths report concern across most risks tested. Concern is particularly high in relation to misuse of personal and financial data (68%), lack of protection if something goes wrong (67%), and concentration of power among large firms (65%). Similar levels of concern are reported for fraud, biased recommendations and financial loss.

Importantly, these concerns extend beyond technical performance. They centre on accountability, incentives and consumer protection. Adoption is constrained more by confidence and perceived risk than by technical capability.

Consumers recognise the potential for AI to deliver benefits in personal finance, but no one benefit dominates

Consumers recognise there is potential for AI to deliver benefits in personal finance. Over half of consumers (55%) already identify at least one potential benefit from using AI for day-to-day financial management, rising to nearly two thirds (63%) among existing AI users.

No individual benefit emerges as dominant. Benefits are spread across multiple capabilities. Among AI users, over a third see value in AI helping them find better-value products or deals (36%), providing more personalised recommendations (35%) and making financial guidance or advice more accessible (35%). Similar proportions also see benefits in supporting faster decisions (32%), improving understanding of financial products (32%) and helping identify potential financial issues earlier (31%).

At the same time, a sizeable group (45%) do not identify any benefits from using AI for day-to-day financial management. This often reflects satisfaction with current approaches, confidence in managing finances independently, or a preference for human interaction.

Safeguards and control are central to adoption decisions

When asked what would increase their willingness to use AI, consumers emphasise safeguards and accountability. The most important drivers are clear protections if something goes wrong (32%) and evidence of accuracy and reliability (30%). Some also highlight the importance of being able to opt out, access human support, and understand how decisions are made.

Convenience is a weaker motivator. Resistance remains significant: 24% of consumers say that nothing would persuade them to use AI in financial services.

There is a gap between perceived authority and actual protection

Consumer understanding of accountability and recourse is limited. When asked what would happen if something went wrong after using a general-purpose AI service for financial advice, only 40% correctly identify that there is no formal route for recourse. The remainder either do not know or believe protections exist.

This creates a risk of unprotected reliance, where consumers may trust AI outputs without understanding the limits of protection. The result is a potential mismatch between perceived authority and actual accountability.

Future adoption is likely to be conditional

Concept testing shows willingness declines as autonomy increases. 36% would use assistive AI, 30% would use AI that acts with permission, and 20% would use fully autonomous AI.

Willingness clearly falls as control is reduced, indicating a preference for oversight. However, the fact that one in five consumers would consider autonomous AI is notable given the novelty of the technology.

Adoption will depend on addressing concerns around trust, control and accountability. While some consumers remain resistant, familiarity and improved safeguards are likely to influence behaviour over time.

Conclusion

AI is now part of everyday life for many and for a growing minority, this is beginning to translate into personal finance. Current usage remains concentrated among those already familiar with AI and is largely focused on assistive applications that reduce complexity without removing user control.

Despite this, over half of consumers already recognise potential benefits. However, confidence remains the key constraint. Concerns focus on data use, reliability and accountability, and many consumers remain reluctant to relinquish control over important financial decisions.

Even among those open to using AI, safeguards matter. Consumers want reassurance around protection and accountability if things go wrong. Currently, only 40% of consumers correctly recognise they have no formal recourse when acting on advice from a general-purpose AI service.

Unlocking further adoption will require clearer accountability, credible safeguards, and stronger consumer understanding of where responsibility lies.

Annex IV - Consumer focus group findings

Eight focus groups, run by FCA qualitative researchers, were carried out in London and Leeds between 10 and 18 March 2026 with a mix of ages, financial confidence and familiarity with technology. Quotas ensured a range of views, including both regular AI users and those with little or no experience. Participants reacted to a series of future-facing examples to explore how they might use AI in financial services.

The findings show how people think about AI in practice. Views were often tentative and sometimes contradictory, reflecting limited experience, uneven understanding and an evolving sense of what AI can do.

Participants described their attitudes to AI through the lens of their broader relationship with money. Differences emerged by life stage, financial resilience, personal responsibility and background. These factors shaped appetite for risk and openness to new technologies.

Consumers were generally open to using AI to support decisions, but this openness was qualified. Across groups, there was a clear preference for AI to assist rather than act independently, reflecting a desire to remain involved in financial decision-making.

"I don't mind suggestions... but at the end of the day, it's my decision."

"I wouldn't give it complete autonomy...in some areas, maybe, but not everything."

Participants most often described AI as a tool for explaining, comparing or organising information. Beyond this, many were unsure how it applied in financial services, limiting confidence and deeper engagement.

Trust was uneven. Consumers were more comfortable relying on AI in lower-stakes situations, such as understanding options, but this reduced where decisions, data sharing or financial outcomes were involved. Attitudes were shaped both by perceptions of AI and by wider scepticism about institutions and incentives.

A further tension emerged around regulation. Many participants assumed safeguards were already in place, despite limited understanding of regulatory boundaries. Money was consistently described as "serious", with protecting hard-earned income carrying strong emotional weight. Alongside financial considerations, some participants raised broader concerns about AI, including its potential effects on jobs and environmental sustainability. These were not always central but contributed to a wider sense of uncertainty about change.

Overall consumer attitudes to AI in financial services are conditional and still evolving. There is openness to AI supporting decisions, but it depends on trust clarity and oversight. Where AI is expected to act autonomously, handle sensitive data or influence outcomes, consumers want clear reassurance on accountability and safeguards.

Annex V - Economic growth impacts

Concept	Potential impact
Automation and Augmentation	Automating or augmenting labour-intensive processes through AI process can help improve operational efficiency, help firms deploy capital more effectively.
Intermediation efficiency thereby reducing costs for consumers.	Bring efficiency within financial services in how firms assess risk, serve customers and process information reducing costs of financial products. This can also help improve customers decision-making and support in quick financial processes.
Productivity Gains	Efficiency already seen within markets and industry e.g. in one Lloyds survey, 59% of institutions reported improved productivity from AI, and responses to our engagement paper also indicated that firms are already realising efficiency gains
Competition Implications	<p>There may also be implications for competition: if AI reduces some service and processing costs, it could weaken the advantage of scale in certain activities and lower barriers to entry, even if trust, data and distribution stay important areas of market power.</p> <p>If consumers and firms rely on a smaller number of models, trained on similar data and producing similar recommendations, decision-making could become more correlated. This could also lead to risks around Systemic herding.</p>
Reduce financial frictions	<p>Lowering the search, switching, administrative and information costs that often lead households and firms to make persistently weak financial decisions. Many poor financial outcomes do not arise because people positively prefer them, but because comparing options, acting and coordinating across products is difficult.</p> <p>In practise there is however potential to leave consumers in expensive debt for too long, holding excess balances in low-yield accounts, or saving too little for retirement</p>
Distributive economic impacts for this channel	<p>AI could expand access to financial services by lowering the cost to serve, improving risk assessment, and making some currently marginal segments more commercially viable. That could support market expansion in areas such as niche lending, while also improving financial inclusion for consumers who are currently underserved or excluded.</p> <p>Fair4AllFinance has estimated annual benefits of around £6.4bn from greater financial inclusion, and around 900,000 people in the UK remain unbanked.</p>
More efficient allocation of household money - where the money moves.	<p>AI could change how savings flow through the economy, because a large stock of household money currently sits in forms that generate very low returns: the FCA has estimated that around £300bn is held in low-interest accounts. If AI tools make it easier for consumers to compare options, act or automate recommendations.</p> <p>If AI makes deposits less sticky, banks may need to rethink pricing, retention incentives and funding models, with implications for how credit is supplied. In some cases, the result may be a more productive allocation of capital; in others, it may create trade-offs for bank funding and SME lending.</p>
Trade-offs for UK Funding	<p>AI could increase the speed, scale and consistency of wider economic effects, e.g. large numbers of customers are nudged towards the same assets or financial choices. The UK-wide gains will depend on adoption areas, but also on how it interacts with funding structures and international competition. If AI-enabled savings and investment flows are directed disproportionately towards international assets, or if UK firms fall behind in adoption and capability, some of the benefits may leak abroad.</p> <p>Equally, if retail funding becomes less stable without stronger alternative channels of finance, there may be trade-offs for domestic lending and investment.</p>

Annex VI - International approaches to AI

International regulatory challenges

Effective international co-ordination is a systemic imperative, but there are associated challenges.

For instance, there is growing divergence among national regulatory approaches (as detailed below), making co-ordination more difficult. Looking ahead to 2030, this tension may intensify as AI becomes more deeply embedded in the financial system and regulatory frameworks continue to evolve in different directions.

Additionally, varying access to AI infrastructure may shape national AI strategies in different ways. AI regulation could increasingly become not only a question of consumer protection, market integrity or operational resilience but also a reflection of industrial strategy, national security and technological sovereignty objectives. These interacting incentives may complicate coordination efforts because the international landscape, rather than being a level playing field, could be shaped by structural interdependencies between jurisdictions with different levels of control over critical AI inputs.

Another co-ordination challenge is on consumer protection. In general, approaches to consumer protection vary significantly across jurisdictions, reflecting different legal traditions, political cultures and risk appetites. Resultingly, there are no emergent interoperable standards on consumer protection. The closest international standards relevant to financial services are the OECD/G20 High-Level Principles on Financial Consumer Protection. There are also various international workstreams on AI and consumer protection. For instance, IOSCO, in collaboration with the FCA, is also progressing work to develop practical tools to support retail investors in identifying and managing AI-related risks, while FinCoNet (the International Financial Consumer Protection Organisation) is set to explore consumer protection impacts of AI, although this work is still at an early stage.

Global standard-setting bodies are already beginning to create useful channels for engagement between financial authorities, firms and parts of the wider AI ecosystem. IOSCO, the FSB and other bodies increasingly seek to engage relevant stakeholders beyond financial services via regular mechanisms such as within the IOSCO Fintech Task Force and through the FSB's ongoing monitoring efforts.

Looking to 2030, the opportunity is to build on these existing channels by considering how engagement can become more regular, forward-looking and ecosystem-wide, supporting earlier assessment of emerging AI-related risks and co-ordinated regulatory responses.

International approaches to AI regulation

Globally, AI regulation is increasingly fragmented with no international consensus anticipated by 2030.

There is broad international agreement on the core principles that should underpin AI governance across sectors. The OECD AI Principles, first adopted in 2019 and updated in 2024, were the first intergovernmental standard on AI and promote trustworthy AI that respects human rights, democratic values, transparency, robustness, safety and accountability. These have been developed into more granular guidance and reporting mechanisms via the Hiroshima AI Process and G20 work. They provide an important global reference point for national AI strategies but remain voluntary and do not create obligations for firms.

However, national approaches to AI regulation are increasingly divergent in the interpretation and implementation of these principles. The main national approaches can be categorised across two dimensions:

Scope: Horizontal, Sectoral or Decentralised

Regulatory style: Prescriptive, Outcomes-led or Voluntary

Meanwhile, Financial standard-setting bodies are also treating AI as a material issue but are not yet seeking to set comprehensive, binding standards. The FSB's 2026 work programme lists AI as a key priority and it has published on AI-related vulnerabilities – identifying third-party dependencies, market correlations and cyber risks as core concerns – as well as on sound practices for AI adoption by financial services firms. IOSCO have also developed a supervisory toolkit on AI as part of its 2026 work programme for use by national financial regulators, providing a useful foundation for supervisory approaches, although ongoing refinement will be important to ensure it remains relevant as AI capabilities and risks continue to develop.

Together, this points to an international regulatory landscape that recognises AI as a growing source of risk, but has not yet reached consensus on whether, or how, to translate that concern into common standards

Examples of national approaches to AI regulation

Jurisdiction	Scope	Regulatory style	Key characteristics	Financial services impacts	Implications and context
United Kingdom	Sectoral	Outcomes based	<ul style="list-style-type: none"> Existing regulators apply current frameworks to AI within remits rather than creating horizontal legislation 	<ul style="list-style-type: none"> FCA applying Consumer Duty, senior managers regime and existing governance frameworks to AI use 	<ul style="list-style-type: none"> Provides flexibility for regulators to apply existing frameworks as AI use evolves, supporting proportionate oversight and adaptability Supportive of innovation Potential need to ensure regulatory clarity
European Union	Horizontal	Prescriptive (for high risk systems)	<ul style="list-style-type: none"> EU AI Act creates cross-sectoral, risk-tiered obligations applying across the single market and extraterritorially 	<ul style="list-style-type: none"> Additional compliance obligations for high-risk use cases (e.g. credit scoring, certain insurance underwriting) layered onto existing financial services rules 	<ul style="list-style-type: none"> Provides a harmonised framework and baseline legal certainty Interaction with existing sectoral obligations creates implementation complexity Seeks to build a trusted single market for AI alongside broader strategic objectives to strengthen digital capability and support strategic autonomy in AI
United States	Decentralised	Hybrid	<ul style="list-style-type: none"> Federal preference for light-touch approach alongside a rapidly evolving variety of state laws 	<ul style="list-style-type: none"> Sectoral regulators generally rely on existing frameworks NIST AI RMF acts as common voluntary reference point 	<ul style="list-style-type: none"> Encouraging of innovation and experimentation Avoids a single rigid framework Differences between state, federal and sectoral approaches may create complexity for firms operating across multiple jurisdictions Overarching national strategy on global leadership in AI
Singapore	Sectoral	Voluntary	<ul style="list-style-type: none"> MAS guidance and toolkits built around FEAT principles through public-private collaboration 	<ul style="list-style-type: none"> Supervisory guidance, lifecycle controls, implementation support and operational toolkits for firms 	<ul style="list-style-type: none"> Strong focus on practical application Collaborative regulatory model based on strong relationships between public authorities, industry and the wider technology ecosystem Reliant on institutional trust and consensus-building Applicability may vary across other jurisdictions with different market structures and institutional arrangements Complements horizontal voluntary frameworks and toolkits developed by the Infocomm Media Development Authority

Jurisdiction	Scope	Regulatory style	Key characteristics	Financial services impacts	Implications and context
Japan	Horizontal	Voluntary	<ul style="list-style-type: none"> • Government promotes non-binding governance principles, guidance and assurance frameworks rather than prescriptive legislation 	<ul style="list-style-type: none"> • Financial regulators primarily rely on existing operational resilience, governance and risk management expectations 	<ul style="list-style-type: none"> • High adaptability • Strong focus on international alignment (Japan led Hiroshima AI process) • Potential need to ensure regulatory clarity
Hong Kong	Combined Horizontal/Sectoral	Hybrid	<ul style="list-style-type: none"> • Regulators issue guidance and expectations within existing sectoral frameworks • Horizontal guidance on data protection and analytics produced by the Office of the Privacy Commissioner for Personal Data (DPCD) and Digital Policy Office (DPO) 	<ul style="list-style-type: none"> • HKMA and SFC emphasise governance, consumer protection and model risk management through supervisory guidance 	<ul style="list-style-type: none"> • Flexible supervisory approach aligned to existing FS regulation • Supportive of safe AI adoption via flexible frameworks and advanced adoption initiatives • Potential need to ensure regulatory clarity
Australia	Horizontal	Voluntary	Existing regulators apply current laws while government has issued voluntary guidelines on AI guardrails	<ul style="list-style-type: none"> • ASIC and APRA generally rely on existing governance, operational resilience and conduct frameworks 	<ul style="list-style-type: none"> • Adaptable and pragmatic approach which supports national strategy to build AI capability • Regulatory direction still evolving with potential future mandatory obligations
China	Hybrid	Prescriptive	<ul style="list-style-type: none"> • Targeted regulations for recommendation algorithms, synthetic media and generative AI with strong central oversight • Government registration process for new AI solutions. • Increasing sector-specific regulation. 	<ul style="list-style-type: none"> • Significant governance and content-control obligations for firms deploying AI systems 	<ul style="list-style-type: none"> • Rapid implementation and strong state coordination with a focus on national security and economic priorities • Strong international engagement on AI governance, particularly in the global south • Need to manage scope for compliance and governance expectations on firms

AI Adoption Initiatives

Globally, regulators are increasingly playing a key role in facilitating adoption, experimentation and capability building.

The UK and Singapore are among the most active jurisdictions in developing financial regulatory initiatives to support responsible AI adoption. The FCA has established internationally recognised innovation services, while Singapore has promoted adoption through public-private collaboration, practical operational guidance and industry toolkits. Other notable jurisdictions include Hong Kong, which has expanded its Gen A.I. Sandbox++ across a range of financial services applications, with supervisory guidance and technical support for firms, and China, which promotes AI adoption through its AI Plus initiative and broader national economic strategy.

Many jurisdictions are explicitly treating AI adoption as important for growth and productivity, rather than focusing purely on compliance and risk mitigation. Regulatory intervention is increasingly being used to build the institutional and market capabilities needed to realise AI's benefits safely.

There is no single international model for encouraging AI adoption. National approaches range from market-led innovation to state-led capability building, with varying sectoral regulatory mechanisms including sandbox testing, supervisory guidance and public-private partnerships. These diverse approaches reflect differences in national priorities, political culture, technological capabilities and resources.

Examples of national approaches to AI adoption

Jurisdiction	Dominant model	FS-specific initiatives	Economy-wide initiatives
United Kingdom	Regulator-led experimentation	<ul style="list-style-type: none"> FCA AI Lab (Sandbox, Live Testing, Sprints, Spotlight, Input Zone) 	<ul style="list-style-type: none"> National AI Strategy – pro-innovation regulatory environment and investment in AI talent, infrastructure and research Regulatory Innovation Office AI Opportunities Action Plan
EU	Compliance-led capability building	<ul style="list-style-type: none"> EU Digital Finance Platform – testing and model training with supervisors 	<ul style="list-style-type: none"> Member state AI sandboxes (under the AI Act) AI Continent Action Plan – compute, data and skills Apply AI Strategy – adoption and sovereign capability Testing facilities under Digital Europe
United States	Market-led / deregulation	<ul style="list-style-type: none"> Limited FS-specific adoption support 	<ul style="list-style-type: none"> National AI Action Plan – sandboxes, AI centres, infrastructure Federal efforts to reduce regulatory burden NAIRR Pilot – AI resources for researchers and educators
China	State-led	<ul style="list-style-type: none"> People’s Bank of China digital finance Action Plan – data governance and AI use 	<ul style="list-style-type: none"> AI Plus initiative – 70% penetration in key sectors by 2027 National standardisation of AI adoption across sectors Infrastructure and open-source ecosystem building District-level sandboxes
Singapore	Public-private partnerships	<ul style="list-style-type: none"> MAS initiatives (Pathfinder, MindForge) MAS-supported AI labs MAS AI grants MAS FinTech Regulatory Sandbox 	<ul style="list-style-type: none"> National AI Strategy AI Verify – testing toolkit Global AI Assurance Sandbox
Hong Kong	Regulator-led experimentation and national capability building	<ul style="list-style-type: none"> HKMA GenAI Sandbox 	<ul style="list-style-type: none"> AI Subsidy Scheme Business guidance Cyberport AI Supercomputing Centre
Japan	Innovation-first strategy	<ul style="list-style-type: none"> FSA discussion paper on AI in finance 	<ul style="list-style-type: none"> AI Promotion Act - promotes private-sector use and investment while avoiding heavy regulation Business guidance Cross-sector sandbox
Australia	Innovation-first strategy	<ul style="list-style-type: none"> ASIC Enhanced Regulatory Sandbox (not AI-specific) 	<ul style="list-style-type: none"> National AI Plan - infrastructure, capability, skills and responsible regulation National AI Centre AI Adopt Programme – funding for SMEs

Annex VII - Infrastructure for agentic finance

The analysis in this annex underpins our recommendation to enable the foundations for agentic finance.

The opportunity

The UK is at an inflection point. It has an opportunity to make decisions now on how AI agents are identified and held accountable in financial services. This will unleash the potential of agentic finance such that by the 2030s, it will be possible for humans to act as observers in a safe and effective way.

As AI becomes embedded across customer journeys and operational workflows, its effectiveness and accuracy will increasingly depend on the underlying infrastructure rather than on individual models alone.

AI-enabled services are expected to operate across multiple firms, products and accounts, but currently prevented by fragmented permissions, limited data interoperability and payments infrastructure designed for human-initiated activity.

Agentic finance in 2030

The Smith School of Enterprise and Environment define agentic finance as finance in which at least one party to the allocation, intermediation, or use of capital is an autonomous AI agent capable of independent economic action.⁸⁸

Respondents to the Mills Review described agentic systems which might be able to 'switch energy, insurance or mortgage providers,' among other activities. By 2030, they may be able to transact, using different forms of money, across retail financial services with delegated authority.

Either: it will accelerate AI by allowing agents to take and execute financial decisions with no human involvement, to optimise recurring bills or provide hyper-personalised financial management, such as autonomous savings allocation or investment. This could benefit consumers from control, protection and access as well as delivering economic benefits through efficiency and competition.

Or: agentic finance could become a significant failure node. If consumers delegate execution authority, the adequacy of consent, authentication, verification and liability arrangements becomes critical.

The building blocks for this form of agentic finance are:

- data
- identity
- authorisation and delegation
- payments
- liability attribution
- supervision and audit at scale

⁸⁸ <https://www.smithschool.ox.ac.uk/research/agentic-finance>

The technology exists and jurisdictions are addressing the questions that underpin agentic finance. The EU is addressing agent authorisation through the AI Act and PSD3. India's Unified Payments Interface (UPI) is integrating with AI assistants at scale.^{89 90} Respondents to the Mills Review agree that key enablers of agentic finance include interoperability standards, digital identity, authentication and real-time payment rails.

The six building blocks that agentic finance requires

1 Data

Data is a critical input to AI in financial services and potentially the greatest enhancer of AI. The Cambridge Centre for Alternative Finance labels data quality as 'the single most significant barrier to scaling AI in financial services.'⁹¹

Data will be fundamental. The IMF identified that data oligopolies allow certain firms with access to superior proprietary data to develop more effective AI models, thereby increasing market concentration.⁹²

Further, these firms consistently outperform competitors, attract more capital, and reinvest in exclusive data sources and advanced AI systems, leading to a self-reinforcing cycle of dominance.

Agents will need three broad categories of data to pursue financial goals effectively:⁹³

- Goal data - this determines what the AI agent is trying to achieve.
- Observation data - what the agent can see about markets, products, prices and the consumer's financial position.
- Outcome data - what actually happened. The record of what the agent did, what resulted and whether it achieved the goal.

Goal data

Goal data is the information layer. It determines what the agent is trying to achieve, why and importantly on whose behalf.⁹⁴ Without high quality, accurate goal data, the agent will pursue the wrong objective competently, which can lead to expensive failures. It is particularly important in retail finance, as consumer goals can have multiple objectives, be dependent on their financial position, and may not be stated explicitly.

Observation data

Observation data is the data that the agent can access and trust about the world when it is going to act. The AI agent needs to understand prices, product terms, counterparty identity, the consumer's state, the agent's own history, the environment's response to its last action. Hui Gong sets out that this is the 'data perception layer.'⁹⁵

89 Alipay's agentic payment infrastructure served 100 million users and processed 120 million agent-initiated transactions in a single week in February 2026.

90 <https://razorpay.com/agentic-payments/>

91 <https://www.jbs.cam.ac.uk/wp-content/uploads/2026/04/ccaf-2026-04-28-global-ai-in-financial-services-report.pdf>

92 <https://www.elibrary.imf.org/view/journals/005/2025/016/article-A001-en.xml?ArticleTabs=Abstract>

93 Importantly, this data must be accurate and representative

94 Identity is discussed in a following section

95 One of the four layers in Financial AI agent architecture <https://arxiv.org/pdf/2603.13942>

Observation data sets out (to the AI agent) what the world looks like. It needs to understand the market and product but also its counterparty (who it is transacting with) and the consumer's financial position, goals, life events and wider context. The AI agent should be able to do this continuously, across many more sources than a human would ever check at machine speed.

A consumer's financial life (and therefore the data related to it) is currently fragmented across numerous providers. Each of these has bespoke product terms, different data schema, different disclosure practices and there is no universal machine-readable product catalogue. Therefore, the observation layer (for the AI agent) is incomplete. Open banking has partially opened it, and open finance aims to open it further.

The FCA is exploring how to strengthen authentication, consent, and trust across open finance, and how this can be underpinned by a robust, interoperable infrastructure.

Open finance could be the observation data layer for agentic finance. It could provide data that is available, portable and standardised across different financial services sectors with clear provenance and traceability to build trust.

Expanded to other sectors (Smart Data) this could provide the building blocks for agents to develop a holistic, real-time picture of individuals.

As Adam Jackson, chief strategy officer at Innovate Finance, states: 'By unlocking high-quality data in a way that secures consumer trust, open finance can be a foundation for widespread adoption of agentic AI'.⁹⁶

There may also be an industry-wide move towards open standards that enable agents to operate across a broader set of financial contexts. The Model Context Protocol (MCP) is one such initiative. MCP is an open standard that gives AI systems the ability to interact with real-world tools and data.

In about a year, MCP has seen growing adoption across industries, including financial services and commerce. Stripe built MCP support to let agents access payment data, create checkout sessions, and manage subscriptions. Square released MCP servers for its payments, catalogue, and customer APIs.

It may be that rather than regulators mandating that banks build standardised open finance APIs, institutions build MCP servers that expose their existing data to agents voluntarily, because doing so makes their own AI products better and their platforms create more value for their customers.

Network effects could drive the enforcement of such a standard. Once enough institutions adopt MCP, those that do not will be disadvantaged as they can't be reached by AI agents.

A risk with a commercially driven standard is that it is only as open as its participants want it to be. There is nothing to stop an incumbent bank from building an MCP server that technically complies but exposes as little as possible.

⁹⁶ <https://www.fca.org.uk/news/press-releases/fca-sets-out-vision-open-finance>

At the moment, MCP and others do not have strong identity guarantees, clear mandate structures and strong audit standards. This is a natural result of them being designed for general use and not for use specifically in regulated finance.

The open finance approach is not without risk. Smart Data (and open finance within it) could act as an accelerant for big tech in financial services.

If big tech companies are the third parties that benefit most from data portability (due to their massive user data, AI capabilities, and embedded distribution) then it may accelerate their penetration into financial services (not enable fintechs). This possibility has been raised in the context of the EU's Framework for Financial Data Access (FIDA negotiations). Credit Mutuel has labelled FID a Trojan Horse for big tech in finance.⁹⁷

Outcome data

This is the information that describes what happened as a result of the AI agent's actions. It provides the feedback that closes the loop between the goal and the reality of what was achieved. This data is also important for supervision of AI agent activity and redress mechanisms.

Data Quality

Firms, consumers and their AI-enabled agents rely on data to make decisions. As set out above, the increased role of AI means data increasingly determines how decisions are formed and executed.

Data quality is therefore foundational. It refers to data that is accurate, complete, timely, consistent and traceable, maintained at a level appropriate to the decision being taken. Where these conditions are not met, outputs become unreliable and increase the risk of poor consumer outcomes. For example, a mortgage application requires verified income data with clear auditability, whereas lower-risk decisions can tolerate less precision.

AI changes how data quality challenges are experienced. It reduces some constraints by enabling the use of unstructured data, with consumers already using AI tools to summarise information and generate suggestions from varied inputs. At the same time, it increases the importance of data quality. In regulated financial services, decisions must be based on trusted, attributable and auditable data. As systems move towards more delegated and autonomous operation, tolerance for error reduces and the need for precision, lineage and governance increases. The challenge shifts from access to data towards maintaining its integrity in operation. This is the system shift on data.

This shift is reinforced by emerging behaviour. Consumers are already uploading financial information and documents to AI services, and in some cases granting ongoing access to their data. Financial data is therefore beginning to move beyond regulated environments, increasing the importance of clear standards for data quality and governance across the ecosystem.

⁹⁷ <https://blog.axway.com/industry-insights/banking-finance/fida-amendments-update>

As set out in the enabling infrastructure section above, AI-enabled financial services will depend on trusted identity, data and execution. AI agents acting on behalf of consumers will rely on goal data, observation data and outcome data. Data quality determines whether this infrastructure can operate effectively.

Financial services will increasingly rely on a combination of structured, programmable data and unstructured data, requiring firms to transform and standardise data before it can be used reliably in AI systems. This makes data collection and transformation processes a core part of data governance. Firms will need to be explicit about which data is required, the format it must take, and the level of quality needed for different decisions. Data must remain accurate and representative in operation, with clear lineage from input to outcome and ongoing monitoring as behaviour, products and markets evolve.

For the FCA, data quality is equally foundational. Effective supervision depends on timely, comparable data that enables cross-firm analysis and the identification of emerging risks. Weak data undermines both firm outcomes and supervisory confidence.

2 Identity

Verification is of central importance as it is a gateway to useful agentic capability.

Firms need to be able to get a verifiable response which sets out who authorised the AI agent and what authority the AI agent has. This is where identity is critical to unlocking the full range of agentic capabilities.

This section thinks about what is needed to be able to hold every AI agent action to the same level of accountability as a human decision. This will require a chain of accountability from the human to the agent, to the action and then back if something goes wrong.

The components are Digital identity to verify that the human is who they say they are and Agent ID - which should exist as a distinct registered entity.

Digital identity

There are many examples of countries using universal digital ID successfully.

In Estonia 99% of citizens have a digital ID. It allows them to vote, sign documents digitally, access healthcare and education, and use banking and online shopping services.

It is estimated that the system saves citizens an average of five working days per year through more efficient processes.⁹⁸ Digital ID and e-signature systems act as a universal Strong Customer Authentication (SCA) which enables more automated processes and complex financial contracts.⁹⁹

98 <https://www.euronews.com/next/2025/09/30/which-countries-in-europe-already-have-a-digital-id-cards>

99 <https://e-estonia.com/solutions/estonian-e-identity/mobile-id/>

In India, the Aadhaar system is an integral part of the Indian tech stack. UPI relies on this digital ID for e-KYC to identify the consumer's identity but also to confirm transactions. The Aadhaar-enabled payment system even resolves a financial inclusion issue by allowing for payments to be made with just an Aadhaar number and a bank account. This removes literacy and device ownership concerns.

The Estonian and Indian models are materially different but demonstrate digital verification's value as a foundational piece of architecture.

Significantly, this utility is a consequence of its near-universal uptake. 99.9% of India's adult population has been issued Aadhaar IDs. Universality makes it useful.

Digital verification in the UK could have benefits which will help to realise agentic finance's potential. These include:

- Continuous verification, which is important where agents may be transacting with numerous firms or other agents across a long time period.
- Regulatory compliance with Money Laundering Regulations.
- Operational efficiencies such as frictionless onboarding.
- Enhanced fraud detection by tying authentication to the human and raising the cost of synthetic identity and impersonation attacks.

We do not propose an approach to digital verification for the UK. Instead, we recognise the utility that universal (or near-universal) digital verification has in enabling agentic finance and propose the need to link it to agent ID.

Agent ID

The UK does not currently define agent identity or require agent verification. There is no registry of AI agent operators. Therefore, if an agent initiates a payment, applies for a product or requests a data release it arrives at a firm with no verified identity, no registered accountability and no way for the firm to distinguish it from an impersonator.

Into this vacuum, two stacks are being developed. One is the agentic commerce and payments stack. The other is the enterprise identity stack. The market is setting the agent verification standard, and the two stacks (and within the stacks) are not interoperable.¹⁰⁰

This fragmented approach risks poor outcomes for consumers, as liability will be governed by the rules of the relevant scheme or platform and revocation is platform-specific and inconsistent.

In a scenario of proprietary standards, incumbents tend to be advantaged as smaller firms face disproportionate compliance costs navigating multiple scheme and platform requirements simultaneously.

A fragmented approach to agent verification limits AI agent capability to what each platform allows.

¹⁰⁰ <https://www.mastercard.com/global/en/news-and-trends/stories/2025/agentic-commerce-framework.html>

Therefore, complex, multi-merchant mandates are not possible or are built on workarounds that lack legal standing. Finally, AI agent activity would be opaque to regulators beyond what the schemes and platforms voluntarily disclose.

3 Authorisation and delegation

A consumer needs to be able to grant their agent bounded, revocable consent to act within defined parameters. From a firm's perspective, they must be able to verify that authority in real-time without requiring customer reauthentication at each step. Without that capability, agentic finance cannot progress beyond the first phases on the spectrum of autonomy.

This is a structural difference to how retail financial services currently assume that a human is the decision-maker at the point of transaction. Current frameworks assume that the customer initiates, the firm authenticates, and the transaction is executed. The open banking consent model requires a human to click 'approve'. Despite being time-bound, scope-bound and revocable, it cannot accommodate an agent acting over days or weeks.

Agentic finance at the highest end of the spectrum of autonomy changes that sequence. The consumer may define their preferences, risk limits and permissions in advance, but the agent may act hours, days or weeks later, and potentially across multiple firms and product types.

The challenge is therefore no longer simply whether the customer consented, but whether the agent's action can be shown to fall within a valid, live and properly scoped delegation of authority.

However, the UK lacks a standardised framework through which a consumer can delegate authority to an AI agent across a wider set of financial actions, with clear limits and accountability.

4 Payments

Payments are the final step in any retail finance journey. An agent can compare mortgages or rebalance a portfolio, but nothing happens for the consumer until money moves.

However, the payments rulebook was written with a human approver (not AI agents) in mind. It is not yet ready for the agent to approve on behalf of someone who may have delegated authority in advance to an AI agent.

The FCA's 2026 Payments Regulatory Priorities report recognises the need to put agentic payments on the policy agenda. Payments are pivotal to amplifying the impact of AI. Without trusted agentic payments, AI in finance remains largely assistive in nature and does not move to the most delegated stages on the spectrum of autonomy.

We are seeing directional forays into agentic payments. Mastercard and Visa are piloting agentic payments and building the tools to prepare for a world in which agents will transact on behalf of consumers and businesses.

Stripe and Open AI have developed the Agentic Commerce Protocol which enables commerce flows between agents and businesses.

For the potential of agentic finance to be unlocked, the payments infrastructure needs to match the capability of AI agents. The agents are there. The development of faster, programmable and interoperable payment infrastructure will be important if agentic payments are to scale.

The Strategy for Future Retail Payments Infrastructure recognises that programmable payments and tokenised forms of money, potentially combined with AI-based innovations, could give consumers and businesses better control over their finances, and lead to customer demand for use cases that do not currently exist.¹⁰¹

Tokenised deposits and stablecoins

We must also discuss the emergence of new digital forms of money.

Tokenised deposits are commercial bank money represented as tokens. They are a digital form of bank deposit recorded on a distributed ledger or programmable platform, rather than held in a traditional account.

Fiat-referenced stablecoins are cryptoassets issued by an organisation seeking to maintain a stable value against an underlying fiat currency, with the aim of functioning as a tokenised form of money on distributed ledger technology. Both developments use distributed ledger technology to offer a payment method that could be faster, cheaper and more secure than traditional payment options.

For AI agents, they enable programmability, which could drive significant innovation in payment use cases and support agentic payments (and agentic finance by extension).

5 Liability attribution

When an AI agent causes financial harm, who is responsible? The answer is important because it directly shapes firm behaviour and impacts the consumer's confidence and willingness to adopt.

A regulated firm that cannot clearly allocate liability for a loss caused by a third-party agent will likely take the rational step of requiring its own human-authenticated confirmation at each step. It will not rely on the consumer's pre-authorisation of the action. A liability framework is important to progress along the spectrum of autonomy. Without it, the consumer experience is full of friction and autonomy is not possible.

6 Supervision and audit at scale

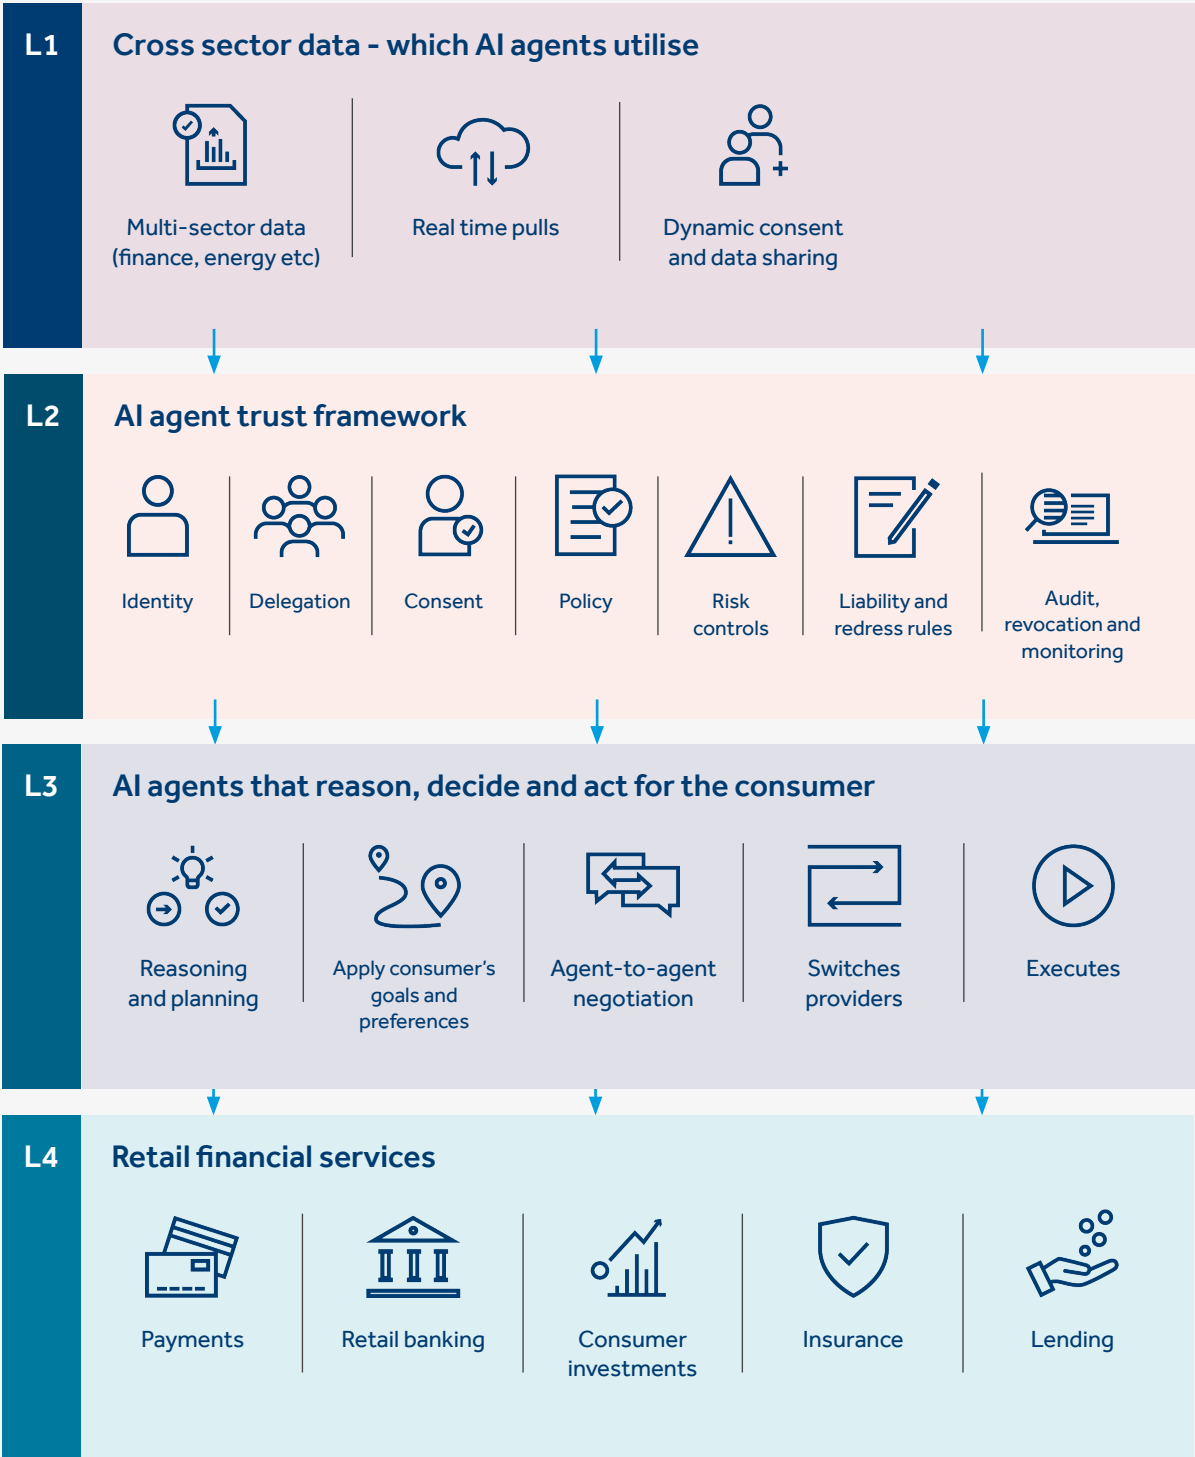
Both regulators and consumers need to understand what an agent did, on what authority, and with what result. At the individual level this can resolve consumer complaints but at the aggregate level, this provides the FCA with a supervisory mechanism to detect systemic patterns.

¹⁰¹ https://assets.publishing.service.gov.uk/media/690de7a247ad122f854627bf/PVDC_Strategy.pdf

How consumers may transact in 2030

The diagram below shows how consumers could use AI agents in 2030 to transact autonomously on their behalf, in an optimal fashion.

This could address a number of longstanding issues: low switching across financial products and excess cash holdings in lieu of investing, as AI agents could overcome customer inertia and optimise for their evolving financial position based on real-time, multi-sector data.





The Mills Review

© Financial Conduct Authority 2026
12 Endeavour Square London E20 1JN
Telephone: +44 (0)20 7066 1000
Website: www.fca.org.uk

All rights reserved