

FCA Review into the long-term impact of AI in retail financial services (The Mills Review)

TheCityUK response

TheCityUK is the industry-led body representing UK-based financial and related professional services. We champion and support the success of the ecosystem, and thereby our members, promoting policies in the UK and internationally that drive competitiveness, support job creation and enable long-term economic growth.

The industry contributes over 12% of the UK's total economic output and employs almost 2.5 million people, with two-thirds of these jobs outside London across the country's regions and nations. It pays more corporation tax than any other sector and is the largest net exporting industry. The industry plays an important role in enabling the transition to net zero and driving economic growth across the wider economy through its provision of capital, investment, professional advice and insurance. It also makes a real difference to people in their daily lives, helping them save for the future, buy a home, invest in a business and manage risk.

Simmons and Simmons LLP have prepared this response on behalf of TheCityUK and synthesised input from several of its members active in UK retail financial services, including major retail banks, retail investment platforms, insurance and wealth management firms. It draws directly on structured interviews and written submissions to reflect the collective views of participating members.

Summary

1. We welcome the Financial Conduct Authority's (FCA) commitment to explore its approach to AI in retail finance through its 'review into the long-term impact of AI on retail financial services' (The Mills Review). The review demonstrates important forward-looking leadership in recognising both the transformative potential of AI and the uncertainty surrounding its future development and impact.
2. A key theme emerging from our consultation with members is the need for targeted clarity on how existing frameworks (such as the Consumer Duty and SM&CR) apply to specific AI use cases, particularly where more complex supply chains and unregulated technology providers are involved. Some members see value in the FCA developing illustrative, non-binding, scenario-based examples in these areas, while others stress that any additional clarification should be firmly grounded in clearly defined, evidenced problem statements and must not harden into de facto prescriptive rules.
3. Taken together, this indicates that further structured engagement between the FCA and industry will be needed to identify the priority use cases and uncertainties where such guidance would be most useful in practice, and to shape its scope and format in a way that preserves the flexibility of the UK's outcomes-based regime.

1. Theme 1: Future evolution of AI technology

- 1.1. *Which emerging or maturing AI technologies do you expect will most transform UK retail financial services from 2030 onwards, and why? Please cite evidence, pilots, or data where possible.*

Looking across the sector, we have identified five clusters of AI technologies that are expected to be most transformative for UK retail financial services from 2030, primarily because they can radically expand access to advice and planning, and reshape how consumers interact with financial services:

(A) **Agentic and conversational AI at the customer interface:** We expect widespread use of AI agents that can scour and aggregate data across accounts and institutions, generate insights and recommendations, and initiate or support customer decisions. These will be embedded in voice and chat channels at scale, re-shaping advice, money management and customer journeys, and potentially disintermediating some current online and adviser-led models.

These tools are seen as transformative because they can materially narrow the UK advice gap. They can provide always-on, low-cost support to consumers who currently do not engage with human advisers, including first-time investors and those who feel intimidated by traditional advice channels. By handling routine queries and pre-work, they can also increase adviser productivity, enabling human advisers to serve more clients and focus on complex or high-stakes decisions.

(B) **AI tools augmenting human-led advice:** Distinct from AI tools at the customer interface, we also highlight tools that support human advisers by automating data gathering, analysis and routine interactions. Plugins that encode professional methodologies (like valuation models and research analysis) could enable banks and fintech apps to offer more sophisticated automated advice to everyday customers. Retail clients may receive insights that previously required professional analysts. These can materially increase adviser productivity while preserving the benefits of human judgement, particularly for vulnerable customers, complex long-term planning and clients who prefer human interaction, ensuring that AI is deployed within existing regulated frameworks.

(C) **Advanced personalisation in retail banking and investment:** AI tools that personalise investment strategies and financial plans at scale are expected to diffuse widely into the retail market, migrating capabilities that are currently more mature in institutional settings. We believe this can democratise access to investment and wealth management solutions across the UK, extending meaningful planning support to consumers who would never pay for, or be offered, traditional advice.

(D) **Agentic AI in payments/commerce:** We anticipate AI agents initiating and executing payments and other transactions on behalf of customers, optimising cashflows across accounts and providers. This “agentic commerce” could significantly reshape retail payment journeys and intermediation, but raises important questions about control, liability and prudential impacts.

(E) **Insurance specific applications:** Members in the insurance sector also highlight AI applications that are transforming retail insurance, including AI-driven claims

automation (e.g. parametric solutions that trigger instant payouts) and AI-supported underwriting tools that enable more personalised, accurately priced policies. Members also point to the potential for disruption in retail distribution of general insurance (GI) products, noting that AI-enabled platforms are beginning to allow consumers to purchase GI products directly through large language models, as seen with examples such as Insurify (US) and Toui (PT) making their apps available via ChatGPT. This development suggests that consumers may increasingly buy GI products directly through AI-enabled channels in the future, further transforming the retail insurance landscape.

1.2. *What do you see as the future potential and direction of agentic AI? What are the implications for retail finance over the coming decade (including accountability, assurance, and market structure)?*

We expect agentic AI to eventually play a central role in retail financial services, particularly in financial planning, payments, savings and investment journeys. With this, we see significant opportunities for the competitiveness of the UK retail financial services market and consumer outcomes. Agentic AI could help close the UK's advice gap by delivering personalised, always-on support and by improving adviser productivity so that human advisers can serve more clients. These tools will democratise access to investment and wealth solutions that have historically been limited to higher-value segments.

However, we identify four main implications:

- (A) **Accountability and liability:** As agents initiate and execute transactions and interact with other agents, firms will need a clear allocation of responsibilities across the AI supply chain (as between the regulated firm and third-party AI providers). Existing frameworks (e.g. Consumer Duty, SM&CR) should ensure that customers continue to benefit from regulated protections even when AI intermediates their decisions.
- (B) **Prudential and market-structure impacts:** Continuous, agent-driven optimisation of customer balances across accounts and products could materially increase the volume and velocity of fund flows, creating new liquidity and stability dynamics that may require firms to augment their approach to risk management.
- (C) **Regulatory Perimeter:** There is concern that agentic AI platforms operated by unregulated firms could provide advice-like services outside the FCA perimeter, amplifying the existing perimeter challenge. We argue that, to protect consumers and ensure a competitive level playing field, the FCA should move towards a “same service, same regulation” approach, supported by clearer expectations or a code of practice for general-purpose and agentic AI providers. However, we recognise that moving towards this approach may necessitate working with HMT where changes to the regulatory perimeter require legislation.
- (D) **Assurance:** We also stress that agentic AI in retail financial services introduces new and novel demands in the context of assurance, explainability and bias management. Decisions and actions taken by agents need to be auditable and traceable, so that firms and regulators can understand and review how outcomes were reached and attribute responsibility appropriately. Control frameworks should include model observability and ongoing monitoring, clear human-in-the-loop escalation points for higher-risk decisions, and guardrails or minimum capability standards for models used in

retail financial services. Consumers should have a high-level understanding of how AI tools work, including the fact that outputs are generated from datasets, models and assumptions, and should receive clear explanations in accessible language rather than technical detail.

- 1.3. What are the UK's comparative advantages and gaps in AI, including compute, data, talent, standards, and regulation relative to other jurisdictions? Which targeted actions might most improve competitiveness?

The UK's existing principles and outcomes-based regulatory approach is viewed as a key comparative advantage versus more prescriptive regimes, such as elements of the EU AI framework, which we fear may be rigid, quickly outdated and potentially inhibit innovation.

We support positioning the UK as a jurisdiction that:

- (A) **Focuses on prioritising consumer outcomes and proportionality over prescriptive and technology-specific requirements.** A supervisory and rule-making approach that focuses on higher risk use cases and evidence-based risk, while allowing lighter touch expectations for lower risk activity, is seen as critical to directing oversight towards material risks and avoiding a chilling effect on innovation across the market.
- (B) **Provides clear, practical expectations** for how existing regimes (e.g. Consumer Duty, SM&CR, operational resilience) apply to AI, including agentic use cases.
- (C) **Maintains UK competitiveness** by ensuring that regulated AI-enabled services remain sufficiently intuitive and convenient, so that consumers choose to use them and benefit from the protections of the regulated sector. We note that if regulated AI journeys are heavily cluttered with risk warnings, pop-ups and mandatory screens, while unregulated AI platforms offer seamless conversational experiences, there is a real risk that consumers may be drawn towards unregulated services and thereby lose the protections of the regulated sector. We advocate for targeted FCA action to calibrate disclosure and interface requirements, so that key risks and protections are still clearly communicated, but in ways that are proportionate and compatible with modern, AI-driven user journeys. This would help to avoid a “dual track-” market in which regulated firms are structurally disadvantaged relative to unregulated AI providers and, most importantly, help ensure that consumers choose to use regulated tools and therefore benefit from the protections afforded to them.

2. Theme 2: Future impact of AI on markets and firms

- 2.1. Could AI systems provide services functionally equivalent to regulated activities such as advice or intermediation, while remaining outside the regulatory perimeter? How might this occur in your market, and what proportion of value could migrate to such unregulated services?

We take the view that AI systems, particularly general-purpose chatbots and agentic platforms, can and do provide advice-like services and outputs that imitate aspects of regulated activities (such as investment advice and guidance) while sitting outside the FCA regulatory perimeter. Consumers increasingly use tools such as large language models to obtain advice-like outputs on investments, mortgages and debt management, often without realising that these tools are not authorised or supervised.

We see several channels through which value may migrate outside the regulated perimeter by 2030:

- (A) **Consumer preference for user experience:** We observe that if FCA-regulated journeys are heavily cluttered with risk warnings, pop-ups and mandatory checks, while unregulated AI tools offer seamless conversational experiences, consumers may be tempted to use unregulated tools instead. This risks shifting advice and guidance activity outside the regulated perimeter, where Consumer Duty protections and suitability checks do not apply. In seeking to create a safe marketplace for AI-enabled retail finance, we encourage the FCA to calibrate the requirements for risk warnings and mandatory information, so that key risks are clearly communicated without making regulated journeys unusably burdensome and unattractive to consumers.
- (B) **Cross-firm agentic intermediation:** We expect that AI agents will increasingly operate across multiple banks and product providers, directing customer funds or transactions between firms. Over time, this could shift value and the primary customer relationship towards agent providers, disintermediating some traditional retail and platform models, and further blurring the distinction between activities inside and outside the FCA perimeter.
- (C) **Anonymity and stigma:** For sensitive issues such as financial distress or bankruptcy, consumers may prefer anonymous AI tools over speaking to a human adviser, further encouraging the use of unregulated platforms.
- (D) **AI-generated financial content and promotions without compliant warnings:** We observe that general-purpose AI tools can and do provide advice-like services and outputs that imitate aspects of regulated activities without pulling through or presenting risk warnings and disclaimers effectively, such as those required under the financial promotions regime. As consumers increasingly encounter AI-generated financial content and prompts on non-FS platforms, these can function as de facto financial promotions and decision-shaping tools outside the scope of existing financial promotions controls, further shifting value and away from regulated advice channels.

From a consumer protection standpoint, this creates risks of unsuitable or misleading advice without Consumer Duty protections, a lack of suitability checks, and inadequate or absent risk warnings. From a competitiveness perspective, we are concerned that the current approach could systematically disadvantage regulated firms relative to unregulated technology providers.

We therefore encourage the FCA to:

- (A) **Move towards a “same service, same regulation” approach**, working with HMT where legislative changes to the regulatory perimeter are required, so that advice-like and intermediation activities face comparable regulatory requirements regardless of provider type.
- (B) **Consider a mix of regulatory and non-regulatory tools**, including consumer education, clearer signposting of what is and is not regulated, warnings about unregulated AI tools, and industry codes of practice, to mitigate harms where direct regulatory coverage is not immediately practical.

3. Theme 3: Future consumer trends

3.1. *How might consumers benefit from AI-enabled retail finance from 2030, and what do you foresee as the greatest risks for consumers?*

We see substantial potential benefits from AI-enabled retail finance by 2030. These include:

- (A) **Hyper-personalised, proactive money management**, with AI analysing income, spending and holdings across accounts to provide timely nudges, recommendations and automated adjustments (e.g. moving funds to avoid overdrafts or optimise savings).
- (B) **Democratisation of access to investment and wealth management tools**, as AI reduces the marginal cost of providing tailored support and enables more consumers to access guidance and portfolio management that were previously uneconomic for human advisers alone.
- (C) **Reduced intimidation and improved financial literacy**, where chat-based AI tools explain concepts in plain language, helping consumers who might otherwise avoid engaging with financial products or advisers.

At the same time, we highlight certain risks:

- (A) **Unsuitable or inaccurate outputs:** Known issues such as hallucinations, opaque reasoning and the use of models that are not sufficiently tailored to UK products, regulation and consumer norms could lead to inappropriate or ill-suited recommendations and poor outcomes. For example, an AI tool developed and trained primarily for a different retail market, where trading and savings behaviours differ, may implicitly assume customer behaviours and risk appetites that do not reflect UK retail customers. Unless such models are carefully adapted for the UK context, they may produce recommendations that are misaligned with UK consumers’ needs, particularly for less sophisticated users who may not recognise these limitations or challenge AI outputs effectively.
- (B) **Over-trust and reduced critical reflection and engagement:** As consumers become accustomed to deferring to AI recommendations and automation takes over more monitoring and optimisation tasks, there is a risk that consumers engage less critically with high-stakes decisions and feel less need to stay actively involved in their financial affairs. We consider that this could create complacency and reinforce already low engagement in areas such as long-term savings and pensions, where consumers should be taking a high level of interest in their decisions and outcomes.

- (C) **Digital exclusion and uneven benefits:** Less digitally engaged customers, or those with thin or fragmented data histories, may receive less personalised or lower-quality support, while digitally confident and mass-affluent customers benefit most.
- (D) **Increased fraud and cyber risks:** AI use lowers the cost and skill needed to mount sophisticated scams and malicious cyber activity (see Question 3.4).
- (E) **Perimeter and Big Tech risks:** Unregulated AI tools provide advice-like outputs without Consumer Duty protections, accountability or adequate risk warnings.

Overall, our members see significant potential for AI-enabled retail finance to improve access, engagement- and outcomes for many consumers. But they also emphasise that these benefits will only be realised if the associated risks are managed proactively through robust firm-level governance and a proportionate regulatory approach. Members stress that the risks outlined above are emerging risks that need to be actively managed, rather than evidence that existing regulatory frameworks or firm oversight are fundamentally inadequate. Firms are already proceeding cautiously in their adoption of AI and developing governance, testing and control frameworks to manage these risks within the current regime. In setting out these risks, members therefore look to the FCA to take an outcomes-based, risk-driven and proportionate approach that supports responsible innovation and the continued evolution of firm-level controls, rather than defaulting to more prescriptive AI-specific rules.

3.2. Which consumer segments might 'win' or 'lose' in this new world of AI-enabled retail finance?

We expect AI-enabled retail finance to deliver earlier and greater benefits for consumers who are already digitally confident and financially engaged, while other groups may face a higher exposure to risk of poor outcomes or exclusion if journeys are not designed inclusively. In responding to this “winners and losers” question, our focus is therefore on which segments are more likely to see improved outcomes, and which are more exposed to the risk of poorer outcomes.

Consumers who are comfortable with digital tools and who already encounter AI in their work or daily lives are more likely to adopt AI-enabled financial services and may see earlier benefits from more convenient, personalised support. This includes digitally confident and mass-affluent customers and, in many cases, younger consumers who are comfortable with AI interfaces. However, members also note that higher baseline trust in AI among some younger users may itself create new vulnerabilities if they are less inclined to question AI-generated recommendations. Over the longer term, if AI-enabled journeys are designed inclusively, there is potential for broader segments - including older, less affluent and late digital adopters - to benefit as well.

Groups that may face a higher risk of poorer outcomes include consumers with thin credit files or fragmented data histories, who may not receive the same level of personalisation or may be misclassified by models; less digitally engaged customers; or those with low financial or digital literacy, who may be excluded from AI-enabled channels or less able to use them safely.

In addition, some members highlight customers with characteristics or drivers of vulnerability as a group requiring particular attention. These customers often rely more heavily on human interaction to ensure their needs are identified and met. They may therefore be at greater risk

if AI-enabled journeys are not designed with appropriate safeguards and clear routes to human support.

3.3. *With the balance shifting between consumer agency and delegation to AI, how might this affect consumer understanding, financial literacy and vulnerability?*

We expect the balance between consumer agency and delegation to AI to shift materially by 2030, with more consumers relying on AI assistants to interpret complex information and make or recommend decisions.

On the positive side, AI can enhance understanding and engagement by:

- (A) Explaining products, risks and options in relatable, plain language and in formats tailored to the consumer, reducing the intimidation that some feel when dealing with human advisers.
- (B) Providing “on demand” explanations and guidance that may improve financial literacy for consumers who would otherwise disengage from financial decisions.

However, we are concerned that increased delegation to AI could:

- (A) Lead to over-reliance and reduced critical thinking, as consumers become used to accepting AI recommendations without fully understanding the underlying trade-offs, especially in high-stakes or complex situations.
- (B) Obscure the distinction between regulated and unregulated advice, undermining consumers’ ability to understand when protections do not apply and when they cannot seek redress when things go wrong.

To manage this within existing consumer protection frameworks, we consider it is important that:

- (A) Consumers should retain meaningful control over whether and how they use AI assistants, particularly in the near term, with straightforward ways to decline or opt out of AI support, and with easy escalation routes to human support.
- (B) Within existing disclosure and consumer-protection requirements, firms should ensure that customers are clearly informed when AI is being used in a journey and what role it is playing (e.g. support vs decision-making).
- (C) Firms should support consumers in achieving at least a high-level understanding of how AI-driven outputs are generated and any key limitations, through accessible explanations rather than technical detail, consistent with existing expectations that information is clear, fair and not misleading.

(D) AI-enabled journeys should be designed so that mandatory regulatory information and key risks are surfaced in ways that AI agents can handle by design, preserving regulatory information flows even when decisions are mediated by AI. We consider that, where unregulated entities provide AI agents that imitate aspects of regulated activities, they should be subject to equivalent expectations so that customers receive the same level of protection regardless of whether they access these services through regulated firms or not.

3.4. *How could AI-driven fraud evolve as consumers increasingly delegate decisions to AI, and what would this mean for consumer agency, harm, and protection in retail financial services?*

We have focused on cyber and financial crime, and we see AI - particularly generative AI - as a step-change in fraud risk by 2030. Key developments include:

(A) **Highly convincing, low-cost scams:** Generative AI dramatically reduces the cost and skill required to produce personalised scam content, including tailored phishing messages, synthetic voices and potentially deepfake video, eroding many of the cues consumers have historically used to identify fraud.

At the same time, we note that AI is also being used defensively, including in biometric and behavioural -analytics-based authentication and fraud detection. These tools can help detect and counter deepfake-enabled scams. But they can also introduce their own risks and the need for robust observability, ongoing monitoring and clear human-in-the-loop escalation points, as well as guardrails or minimum capability standards for models used in fraud controls.

(B) **More sophisticated malicious code and attacks:** AI tools can be used to design malware and more complex attack patterns, raising the bar for firms' cyber-defence capabilities.

We expect that heavy reliance on repeated in-app warnings and friction will not be sustainable as AI-driven fraud evolves. Instead, we anticipate that firms' approach to managing financial crime risk will need to continue to evolve towards, and be characterised by:

(A) **Layered, adaptive controls**, including behavioural analytics and behavioural biometrics, which many organisations are already using for fraud detection and prevention, together with out-of-band contact (e.g. reaching customers through alternative channels when behaviour looks unusual) and dynamic authentication. We also highlight the importance of non-financial agentic providers participating in established fraud-prevention information-sharing arrangements (for example, through industry bodies such as CIFAS), so that AI-enabled fraud patterns can be identified and tackled across the wider ecosystem.

(B) **Enhanced fraud and cyber frameworks** that are recalibrated to reflect AI-enabled threats. This includes enhanced observability and ongoing monitoring of AI models used in fraud controls, strengthened data security around model training and deployment, and robust incident detection and response capabilities. Members emphasise that these enhancements can be achieved within the existing outcomes-based regulatory frameworks, and that additional requirements focused specifically on managing cyber and resilience impacts in the AI context are not necessary. We encourage the FCA to recognise the extent to which current fraud, cyber and operational-resilience frameworks already provide a foundation for managing AI-enabled threats, and to focus on proportionate, risk-based expectations with targeted additions where genuinely novel risks arise.

We also caution that it is unrealistic to expect consumers to verify firms or investments by manually checking FCA registers or websites, especially as decisions are increasingly mediated by AI agents. One suggestion is that, over time, there may be value in the FCA exploring options for making certain core elements of its register of information and consumer warnings available in content-managed, machine-readable formats that firms could integrate into AI agents, so that consistent regulatory messages can be surfaced in real time to consumers. However, we recognise that any such approach would require careful consideration of feasibility, oversight and how responsibilities are allocated across technology providers and regulated firms, potentially alongside wider cross-sector initiatives.

3.5. *What might help make AI-driven decisions more understandable and trusted by customers, including how the use of AI may be monetised?*

We agree that trust is pivotal to consumer adoption of AI-enabled retail finance, and that it cannot be assumed that trust will automatically follow product availability.

To build and sustain trust while supporting UK competitiveness, we highlight several priorities:

- (A) **Transparency and disclosure:** Consumers should be clearly informed when AI plays a material role in delivering a service or generating outputs, particularly where this affects whether they are receiving regulated advice or unregulated outputs. Clear disclosure in these contexts is seen as a key enabler of trust.
- (B) **Human escalation and boundaries:** We see clear routes for consumers to escalate complex or disputed matters, as part of a broader set of safeguards. Human review should complement, rather than replace, other controls and deterministic safeguards built into AI-enabled journeys.
- (C) **Clear monetisation and incentives:** Any commercial incentives behind AI outputs (e.g. partner-funded recommendations) should be clearly disclosed.
- (D) **Robust governance and accountability:** We consider it important that AI-related risks are integrated into existing governance and risk-management frameworks, with clear accountability under SM&CR and the use of auditability and explainability to support both internal oversight and external supervision. This is seen as a way of sustaining trust over time while building on, rather than duplicating, existing governance structures.

- (E) **Regulatory labelling and standards:** Clear labelling and standards could help consumers distinguish regulated, trustworthy AI tools. In practice, this may be more likely to develop through industry-led standards and voluntary labelling schemes, potentially supported by FCA guidance, rather than through FCA-mandated kitemarks or watermarks for all AI-generated content, particularly where non-regulated entities are involved.

4. Theme 4: Future regulatory approach

4.1. What are the opportunities and challenges for the FCA in ensuring an outcomes-based approach to retail regulation in an AI-enabled FS industry?

The UK's existing principles and outcomes-based approach to AI-related regulation in financial services is an important strength that helps position the UK as an innovation-friendly jurisdiction. We explicitly contrast this favourably with more prescriptive regimes, such as elements of the EU framework, which are comparatively rigid and at risk of inhibiting innovation and reducing jurisdictional competitiveness. There is therefore a clear opportunity for the FCA to reinforce and build on this model and thereby position the UK as a leading innovation-friendly jurisdiction for AI-enabled retail finance.

An outcomes-based approach gives the FCA flexibility to respond to actual, evidence-based risks as AI use cases evolve, without unnecessarily constraining innovation through detailed, technology-specific rules. We emphasise that further intervention should be risk-driven and proportionate, with a preference in many areas for "best practice" support and guidance over hard regulatory requirements, particularly where risks are still emerging. This is central to maintaining the UK's competitiveness while enabling the safe deployment of AI in retail finance.

While the existing outcomes-based framework remains appropriate, some members would find it helpful for the FCA to provide targeted, illustrative scenario-based examples of how it interprets existing obligations (including Consumer Duty and SM&CR) in specific AI use cases where uncertainties arise. However, other members caution that even illustrative examples can, in practice, become de facto prescriptive rules, which would not reflect the intended flexibility of the outcomes-based approach. They therefore stress that any guidance should be grounded in clearly defined problem statements and carefully framed, to avoid hardening into prescriptive requirements and to preserve the flexibility of the current framework.

A key challenge for the FCA will be to provide clear, consistent supervisory expectations that support iterative improvement in firms' use of AI, while maintaining strong accountability for consumer outcomes and market integrity. Members consider that a risk-based approach that focuses on governance, effective risk identification and remediation, rather than on technology-specific rules, would best support responsible innovation.

We stress that the FCA's focus should be on overall risk-benefit, governance, remediation and learning, rather than expecting zero defects. If the FCA can strike this balance, we consider that the UK can strengthen its position as a competitive jurisdiction that attracts AI-related investment, while maintaining high standards of consumer protection. We also note that firms are already integrating AI into existing governance and risk frameworks and see value in the FCA recognising this work when calibrating expectations and supervisory approaches.

4.2. Are the key FS regulatory levers (Consumer Duty, Operational Resilience, SM&CR, Critical Third Party (CTP) regime, etc.) suitable to manage future risks and to enable firms to fully take advantage of AI?

We generally consider that the existing suite of FS regulatory levers provides a workable framework for managing AI-related risks while enabling innovation, provided they are applied and - where necessary - clarified in a way that reflects AI-specific use cases.

- (A) **Consumer Duty:** We believe the Consumer Duty can work for AI but note that many AI-enabled journeys will involve more complex supply chains and unregulated technology providers, rather than the predominantly 1:1 provider-customer relationship for which the Duty was designed. In that context, some members would welcome greater certainty over how the Duty applies where services are delivered using unregulated technology providers, including clarification on liability allocation, auditability and consumer protections. Others stress that any additional guidance should be linked to specific, clearly articulated uncertainties, to avoid turning illustrative examples into de facto prescriptive rules.
- (B) **SM&CR:** It is important to clarify how responsibility for AI models and use cases should be allocated within SM&CR, ensuring that there is clear senior-manager accountability for AI governance, model risk and consumer outcomes.
- (C) **Operational resilience and critical third parties (CTP):** The CTP regime represents an important mechanism for addressing systemic risks arising from critical third-party dependencies. In principle, it could be used, where appropriate, to cover AI providers whose failure or disruption could pose a systemic risk to the financial sector.
- (D) **Use of existing consumer-protection frameworks:** We emphasise that aspects of the Consumer Duty and related frameworks can be applied or extended to AI activities, avoiding the need for a wholly separate AI regime. Some members would find it useful for the FCA to continue engaging with industry and, where specific uncertainties arise, to provide illustrative, non-binding examples to assist with implementation. Other members stress that any such guidance should remain problem-driven to avoid hardening into prescriptive rules, precisely because overly specific examples could harden into checklists and limit firms' flexibility and innovation, undermining the flexibility of the outcomes-based regime.

4.3. Do you have views on the way the FCA should improve or develop its approach to supervision and/or enforcement to respond to increased AI use in the future, including using AI itself?

We encourage the FCA to adopt a constructive, adaptive supervisory approach to AI that supports innovation while maintaining high standards of consumer protection.

We highlight several priorities:

- (A) **Risk-based supervision and tolerance of inevitable error:** AI models, like other complex systems, will not be error-free. It is therefore important that supervision focuses on overall risk-benefit balance, governance and effective remediation, rather than on technology-specific rules. A risk-based, outcomes-focused supervisory approach is

better suited to supporting responsible innovation while maintaining strong accountability for consumer outcomes and market integrity.

- (B) **Alignment between supervisory practice and the FCA's strategic aims:** Day-to-day supervision and enforcement should be clearly aligned with the FCA's stated aims of encouraging safe AI adoption, improving consumer outcomes under the Consumer Duty, and supporting competition, growth and UK competitiveness. The supervisory approach should remain aligned with the FCA's public commitment to support innovation and avoid becoming disjointed from those aims in practice. In practice, this means supervisors balancing vigilance over emerging harms with a clear recognition that responsible experimentation within the regulatory perimeter is necessary if AI is to deliver the benefits to consumers and to the UK that the FCA itself has highlighted.
- (C) **Supply-chain and third-party oversight:** As AI supply chains become more complex, regulated firms may face information asymmetries and limited leverage with some major technology providers. They consider it important that the FCA takes these practical constraints into account when supervising how firms manage AI-related third-party risks under existing outsourcing and operational-resilience frameworks, and that any evolution of the framework remains proportionate and aligned with existing principles-based approaches.
- (D) **Use of AI by the FCA itself:** We see opportunities for the FCA to use AI tools to make better use of existing data (for example, to identify patterns in firms' regulatory returns) without increasing reporting burdens on firms. Any FCA use of AI in supervision and enforcement should support, rather than replace, human-led decision-making, so that regulatory judgements continue to reflect context and nuance.

A supervisory stance which is predictable, proportionate, and clearly aligned with the FCA's strategic objectives on AI and competitiveness will be critical to encouraging ongoing investment in AI within the regulated perimeter. This reflects experience in other areas where shifts in regulatory approach have materially influenced firms' appetite to invest.

4.4. *In what ways can the FCA continue to support growth and competitiveness in an AI-driven financial services industry in the future?*

We see significant opportunities for AI to improve productivity, expand access to financial services and enhance the UK's position as a global financial centre, provided the regulatory environment remains innovation-friendly, proportionate and predictable.

Key ways in which the FCA can support growth and competitiveness include:

- (A) **Maintaining a distinct, outcomes-based approach.** We strongly favour the UK's principles-driven model over more prescriptive regimes, viewing it as better suited to fast-evolving AI technologies. We urge the FCA to resist adopting rigid, technology-specific rules that could quickly become outdated and deter investment.

- (B) **Reducing uncertainty through clear, practical guidance.** Where specific uncertainties arise, some members would find it helpful for the FCA to continue engaging with industry and, where appropriate, to provide illustrative, non-binding examples of how existing requirements apply to AI use cases. Other members are cautious about seeking additional guidance unless there are clearly defined problems to solve and emphasise the importance of preserving the flexibility of the current outcomes-based regime and avoiding illustrative material hardening into de facto prescriptive rules.
- (C) **Ensuring a level competitive playing field with unregulated technology firms.** Regulated firms should not be structurally disadvantaged relative to unregulated AI providers because of how existing requirements are implemented - for example, through disproportionate friction in regulated journeys compared with unregulated tools. We therefore encourage the FCA, when calibrating disclosure and interface expectations, to consider the risk that excessive friction could push some consumers towards unregulated services and away from the protections of the regulated sector. The underlying perimeter and conduct regimes are broadly fit for purpose. More detailed perimeter issues are addressed elsewhere in this submission.
- (D) **Supporting common standards and infrastructure.** The FCA can play a useful convening role in developing baseline capability standards for AI models used in retail financial services, and in exploring, over time, ways in which certain core elements of regulatory expectations could be made available in formats that are easier for firms to embed into AI systems. Any such work would need to take account of feasibility, oversight and the respective responsibilities of regulated firms and technology providers.

4.5. *Are there other regulatory frameworks (UK or international, other non-FS sectors) which the FCA might consider or emulate to respond to increased AI use in retail financial services?*

Several domestic and international frameworks could inform the FCA's approach to AI in retail financial services:

- (A) **Singapore's "living frameworks":** Singapore's regulator issues regularly updated policy notes and guidance that build on existing international standards to maintain a shared understanding between regulators and industry without creating rigid new obligations. A similar approach in the UK could allow AI-related expectations to evolve with technology while preserving the advantages of principles-based regulation.
- (B) **ICO guidance:** The Information Commissioner's Office scenario-based guidance is a model the FCA could use to clarify expectations for AI use under the Consumer Duty and other regimes, providing concrete examples without overly prescriptive rules.
- (C) **Critical third-party and operational resilience regimes:** The CTP regime is designed to address systemic risks arising from disruption at critical third-party AI providers. Where AI providers pose comparable systemic risks, they could fall within the scope of the regime. The FCA has already indicated that it may be appropriate to extend the framework to AI providers, which might necessitate adjustments to the framework to ensure it appropriately captures systemic risks arising from AI-related dependencies.

(D) **EU AI transparency and watermarking rules:** We reference the EU AI Act's watermarking and transparency provisions as an example of how difficult it can be to operationalise such requirements in practice. We see this more as a cautionary example than a model to emulate in the UK, especially given the UK's preference for principles and outcomes-based regulation.

Overall, we favour a UK approach that combines the flexibility of Singapore-style living frameworks, the concreteness of ICO-style scenarios, and selective learning from existing FS tools (e.g. CTP and operational resilience), rather than importing any single regime directly.

We appreciate the opportunity to respond to this review and would be happy to answer any follow-up questions.

End of document.