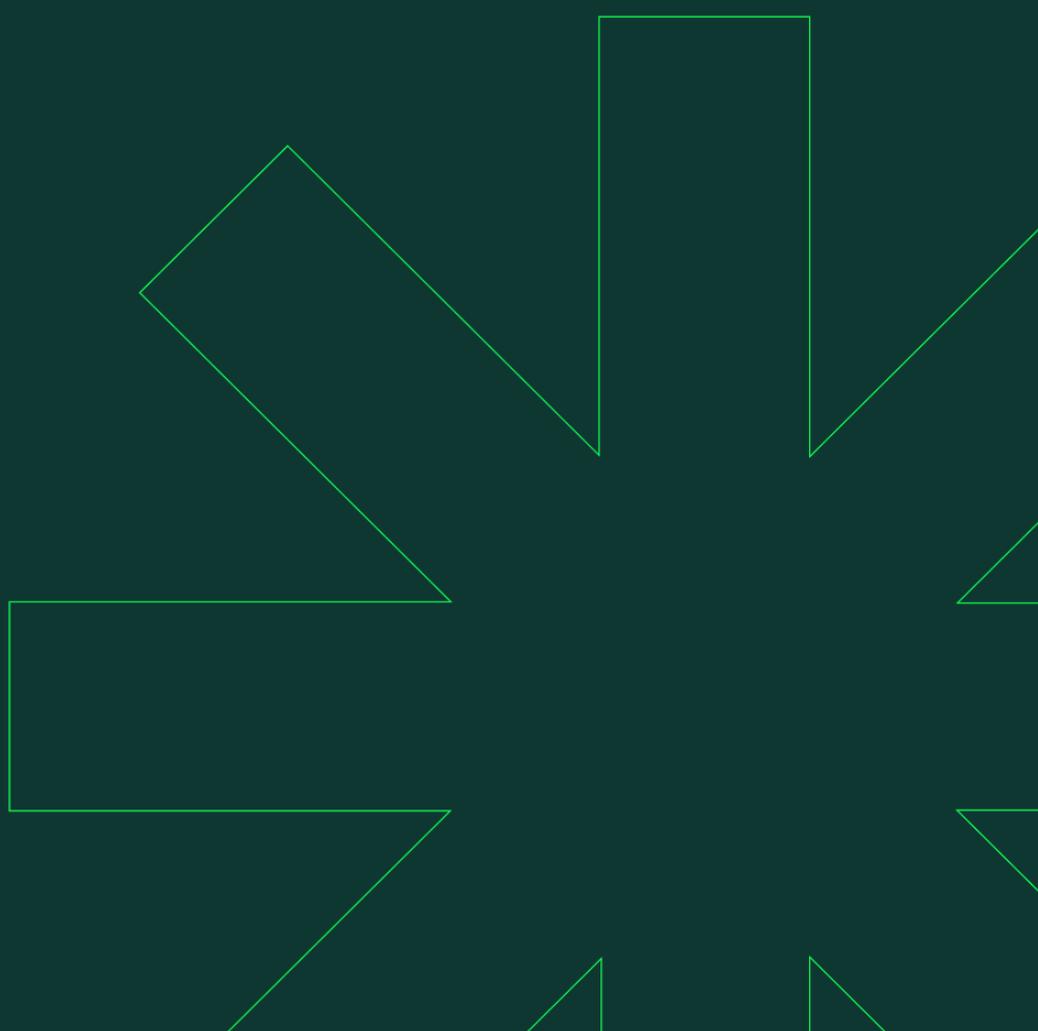


# AI Vision for Fintech

Startup Coalition's Response  
to the Mills Review

February 2026

STARTUP  
C\*ALITION

The bottom right corner of the page features several overlapping, hollow geometric shapes in a light green color. These shapes include rectangles, triangles, and a diamond, all oriented in various directions, creating a modern, abstract graphic design.

# About Startup Coalition

Startup Coalition is the policy voice of UK tech startups and scaleups in Westminster. Since 2010, we have worked to engage on behalf of tech startups in public policy debates in the UK across a range of critical priority issues including access to finance, immigration and skills, and technology regulation.

We fight for a policy environment that enables early-stage British tech companies to grow, scale and compete globally. We have over 4,000 startups in our network and have been instrumental in building proactive coalitions of businesses and investors on issues that are integral to the health of the UK's startup ecosystem. Startup Coalition works directly with the Government across a range of issues. We represent the startup community on the Smart Data Council.

# General Remarks

Startup Coalition welcomes the Mills Review as a timely recognition that the UK is approaching a genuine inflection point in retail financial services: not just because “AI” is improving, but because AI is becoming more autonomous, embedded in infrastructure, and capable of acting on behalf of consumers. It is also radically improving financial products and services.

Our view is that the future described in the Review is already here: a joint survey of UK financial firms found that 75% already used AI, with a further 10% planning to do so within three years, and that over half of AI use cases involve at least some automated decision-making. The policy question, therefore, is not whether to “allow” AI; it is whether the UK’s regulatory and data infrastructure will channel AI into outcomes that are pro-consumer and pro-competition, and maintain the UK’s place as the home of Fintech innovation.

Having led the world in the last wave of Fintech innovation, the intuitive interfaces trailblazed by Monzo, Revolut and Starling are now table stakes. The FCA cannot predict what the next assumed features will be in financial services, but we must expect them to be AI enabled.

Our core proposition is straightforward: the FCA should remain outcomes-based and risk-proportionate, but should be more deliberately pro-innovation in the enabling layer, especially Open Finance, payments rails, and test-and-learn regulatory pathways, because these are the prerequisites for ambitious, responsible “agentic finance” and the future of fintech innovation in the UK.

**If the FCA leans into this new era, the UK stands to benefit enormously.**

# Consultation Response

## Theme: Future Evolution of AI Technology

**Question: AI technologies — Which emerging or maturing AI technologies do you expect will most transform UK retail financial services from 2030 onwards, and why?**

From 2030 onwards, the most transformative technological shift in financial services will be from “AI as a feature” to “AI as an operating layer” across the consumer’s financial life: systems that can reason over multi-product financial positions, recommend actions, and execute them within defined guardrails.

However, the key point for policy is that “AI technologies” are not only models. They include the full AI value chain and the financial rails that determine whether AI delivers broad-based consumer benefit. Indeed, AI has already significantly improved conventional financial services in 2026.

A practical taxonomy for retail finance beyond 2030 includes:

### **1. Agentic and autonomous systems (including tool-using agents)**

Agentic AI systems will increasingly shift from chat-based explanation to action-based orchestration (e.g. analysing transactions, constructing plans, switching products, initiating payments, and continuously monitoring affordability). The Review foregrounds this shift, and the Bank of England and FCA’s 2024 survey of 118 regulated financial firms found that semi-autonomous decisioning is already common while fully autonomous use remains limited, suggesting a regulatory opportunity to shape the next phase rather than merely react to it.<sup>1</sup>

### **2. Algorithms everywhere: decisioning, pricing, underwriting and market conduct.**

Even before fully autonomous consumer agents, “classical” ML and optimisation algorithms will drive continuous underwriting, dynamic pricing, fraud detection, and personalised servicing. The systemic risk question is not whether algorithms will decide, it is whether the surrounding governance ensures these decisions are contestable, explainable and aligned with consumer outcomes at scale.

### **3. Foundation and multimodal models as the interface layer for finance.**

Foundation models (including large language models) are rapidly becoming part of financial services workflows: the Bank of England’s survey reported that foundation models already constitute a meaningful share of AI use cases, consistent with the broader shift toward natural-language interaction and automation.

---

<sup>1</sup> <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

#### 4. Adversarial AI and AI-enabled fraud as a parallel technology curve.

The evolution of AI in retail finance is inseparable from the evolution of attacks on AI and the use of AI by malicious actors. The Financial Stability Board has highlighted cyber risk and the potential for GenAI to accelerate fraud and disinformation, while the BIS has emphasised that cybersecurity risks and adversarial attacks are a key challenge when sensitive data and algorithms are shared, including with third parties.<sup>2</sup>

This matters directly for retail markets: AI-generated disinformation can move faster than traditional crisis management. Evidence cited by Reuters from a UK study suggested a large share of customers would consider moving money after exposure to AI-generated falsehoods, and that small paid amplification could plausibly trigger large deposit movements, an operational and communications resilience challenge in an era of digital banking.<sup>3</sup>

#### 5. AI assurance platforms as a core enabling technology.

The Review anticipates AI assurance platforms, tools to monitor, audit and evaluate AI systems, as “critical for trust and resilience” in the future landscape. This is not a secondary issue: it is the mechanism that can make ambitious AI compatible with outcomes-based regulation by turning model risk into measurable, monitorable controls.

#### 6. AI in a quantum age: two distinct issues.

First, quantum computing is a potential accelerator for optimisation, machine learning and stochastic modelling in finance, and the FCA’s own research has concluded that near-term commercial applications may be viable and that regulators can support innovation while managing risk, explicitly suggesting new regulations are unlikely to be required in the near term, provided existing themes (explainability, fairness, operational resilience) are applied intelligently.<sup>4</sup>

Second, quantum capability changes the security baseline: the Bank of England has warned that preparations need to happen now for a post-quantum future and that quantum computing could undermine the asymmetric cryptography underpinning the financial system, with “harvest now, decrypt later” risks.<sup>5</sup>

Internationally, the National Institute of Standards and Technology has finalised initial post-quantum encryption standards, and UK guidance has begun to focus organisations on migration planning, both highly relevant to financial services’ long-lived confidentiality requirements (e.g. identity data, customer records, trading and payments infrastructure).<sup>6</sup>

<sup>2</sup> <https://www.fsb.org/uploads/P14112024.pdf>

<sup>3</sup>

[https://www.reuters.com/technology/artificial-intelligence/ai-generated-content-raises-risks-more-bank-runs-uk-study-shows-2025-02-14/?utm\\_source=chatgpt.com](https://www.reuters.com/technology/artificial-intelligence/ai-generated-content-raises-risks-more-bank-runs-uk-study-shows-2025-02-14/?utm_source=chatgpt.com)

<sup>4</sup>

<https://www.fca.org.uk/publication/research-notes/quantum-computing-applications-financial-services.pdf>

<sup>5</sup>

<https://www.bankofengland.co.uk/report/2025/the-boes-approach-to-innovation-in-ai-dlt-quantum-computing>

<sup>6</sup>

The strategic implication is that the FCA should treat “AI technologies” as an ecosystem question: models + data + compute + assurance + cyber + execution rails. The regulator can be pro-innovation by targeting its interventions at the enabling layer (portability, interoperability, testing infrastructure, assurance expectations).

## **Theme: Future impact of AI on markets and firms**

### **Question: Market structure and customer passthrough - How might AI change concentration in your market?**

AI can either intensify concentration or significantly lower barriers to entry; the outcome hinges on who controls (a) data access, (b) distribution interfaces, and (c) critical AI inputs such as cloud and model providers. The Review flags “winner takes most” dynamics (data feedback loops, economies of scale, network effects) as a key uncertainty for market structure.

Evidence already points to concentration pressure in the AI supply chain. The Bank of England’s survey reported that around a third of AI use cases are third-party implementations and that the top three providers account for very large shares of cloud, model and data providers reported by firms, implying both operational dependency and potential competitive bottlenecks.<sup>7</sup> Global assessments reinforce this: the Financial Stability Board has highlighted third-party dependencies and service provider concentration as a key vulnerability that can increase systemic risk.<sup>8</sup>

At the same time, AI can empower challengers by compressing unit costs for onboarding, servicing, compliance and customer support, particularly where firms can access high-quality data and interoperable rails. The Bank of England’s survey reported that firms expect the biggest growth in benefits over the next three years to be operational efficiency, productivity and cost base.<sup>9</sup> This is consistent with broader evidence that generative AI assistance can materially raise customer support productivity, especially for less experienced workers, suggesting that smaller firms can scale service quality faster than before if they can access the right tools and data.<sup>10</sup>

Cost reduction does not automatically translate into customer benefit, however. Whether savings are passed through as lower prices depends on competitive intensity and switching costs, factors that AI itself can reshape. A pro-competitive scenario depends on data portability and reliable execution: if incumbents can degrade rails, introduce friction, or restrict access, AI becomes a reinforcing technology for existing market power rather than a disruptive force.

---

<https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>

<sup>7</sup> <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

<sup>8</sup> <https://www.fsb.org/uploads/P14112024.pdf>

<sup>9</sup> <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

<sup>10</sup> [https://www.nber.org/papers/w31161?utm\\_source=chatgpt.com](https://www.nber.org/papers/w31161?utm_source=chatgpt.com)

**Question: Self-reinforcing dynamics - What evidence do you see of ‘winner takes most’ dynamics?**

The Bank for International Settlements has previously argued that AI can increase systemic risk through “data uniformity”, “model herding”, and network interconnectedness, and that reliance on a small number of data and model providers increases the risk of procyclicality and correlated failures.<sup>11</sup>

This has both competition and stability implications: if many consumer finance providers converge on the same few models and datasets, the market may converge not only on suppliers but on decisions, pricing, credit approvals, and portfolio shifts, reducing diversity and increasing correlated outcomes.

**Question: Control of the customer relationship - Who will control the primary customer relationship by 2030 onwards?**

We expect the “primary customer relationship” to become interface-mediated: a consumer’s main interaction may be with an AI layer (a bank’s own assistant, a specialist intermediary, a device-level agent, or a platform agent), rather than with product providers directly. The Review explicitly anticipates that value may shift to those who control AI-enabled interfaces and that some AI-enabled services could capture value while remaining outside the regulatory perimeter.

This creates a fork in the road: on the one hand, a pro-competition outcome is plausible if consumers can choose and switch their AI intermediary easily, and if their agent can access and act on comprehensive financial data across products. The UK’s emerging Smart Data architecture is designed to extend Open Banking-like sharing to a wider range of products in financial services (such as savings, investments, pensions and insurance) and beyond, and to enable new services and stronger consumer control. Startup Coalition sits on the Smart Data Council and supports the use of Smart Data powers under the Data Use and Access Act 2025, where the market is failing to deliver on consumer demand for data access and portability.<sup>12</sup>

For “reading” data, we believe that Smart Data schemes are aspirational across financial services, the energy market, property, retail and transport.<sup>13</sup> In these instances, the in-scope data should be all that belongs to the consumer under GDPR, data must be made available for free or at nominal cost, in real-time, and via standardised APIs. It is also aspirational for each sectoral scheme to be interoperable with one another, and for there to be cross-regulatory streamlining of permissions. For instance, a regulated Account Information Service Provider under the Open Banking regulations should be able to connect to the APIs for every single other regulated scheme, since no other data sets will be as sensitive as financial datasets.

However, a more concentrated outcome is plausible if a small number of platforms control the interface, the data, and the execution rails, creating new lock-in and “pay-to-play” distribution in which financial firms compete for the attention of the AI rather than the human.

<sup>11</sup> <https://www.bis.org/publ/work1194.pdf>

<sup>12</sup> <https://startupcoalition.io/news/making-smart-data-happen/>

<sup>13</sup> [https://api.startupcoalition.io/u/2025/12/Startup-Coalition\\_License-to-Share\\_FOR-RELEASE.pdf](https://api.startupcoalition.io/u/2025/12/Startup-Coalition_License-to-Share_FOR-RELEASE.pdf)

For the FCA, the central competition challenge is interoperability. The Bank of England's survey already showed many firms have only a partial understanding of the AI they use (often because of third-party models), which increases both operational dependence and the risk that the interface layer becomes a de facto gatekeeper for market access.<sup>14</sup>

Finally, there is a challenge with "write" access APIs, whereby users can consent to third parties editing data or initiating action via an authorised third party (ATP). Today, the only existing regulated "write" access APIs are payment initiation APIs under Open Banking, or its derivatives, such as sweeping or variable recurring payments. Startup Coalition believes that incumbent opposition to Open Banking and similar Smart Data schemes comes from fear of increased competition, and disintermediation that occurs when the user can interact with the service they provide *without interacting with their proprietary interface* (namely through a regulated ATP). This is a poor argument against Smart Data schemes; however, we are alert to the broader consequences of "write" access APIs proliferating across the economy.

Combined with agentic AI, the rapid increase in automated actions across the economy via interoperable APIs could be seismic, and the Government should rapidly undertake an examination of this scenario to understand the extent of the impact. For instance, we have heard evidence from the building society sector that rapid, automated movement of money from savings accounts could undermine mortgages, and therefore the property market. It is vital that the Government assess where such risks exist, not necessarily to avoid taking action to introduce these schemes, but to do so with their eyes wide open. Indeed, identifying these weaknesses could support the government in identifying where specific sectors are actually currently too vulnerable, weak, or unsophisticated. To take the building society example further: if agentic AI at scale is an integrity risk to building societies, then perhaps the building society sector is not fit for purpose in 2026. In this way, AI can be helpful in identifying archaic business models that require reform.

**Question: Regulatory perimeter - Could AI systems provide services functionally equivalent to regulated activities while remaining outside the perimeter?**

Yes; and evidence suggests this is already beginning. A major consumer study found that 51% of UK adults use AI search tools and that a meaningful minority report using AI "often or always" instead of finance professionals, while trust in these systems remains relatively high among users.<sup>15</sup> It is clear that the cat is out of the bag, and that a majority of consumers now demand these services.

At the same time, Which? testing found that popular AI tools are inconsistent in output according to their in-house experts.<sup>16</sup> Meanwhile, recent Sky News analysis in partnership with Hargreaves Lansdown criticised the investment advice given by ChatGPT, Microsoft Co-pilot and Google Gemini for reasons including US centricity, a lack of diversification, and

---

<sup>14</sup> <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

<sup>15</sup> <https://www.which.co.uk/policy-and-insight/article/consumer-use-and-attitudes-towards-ai-search-tools-aTnr81n3FOQI>

<sup>16</sup> <https://www.which.co.uk/policy-and-insight/article/chatgpt-and-gemini-among-ai-tools-giving-risky-consumer-advice-which-finds-aBnBP0I2CE0T>

the suggestion to invest in real estate.<sup>17</sup> Indeed, in our own recent report, All About the Gains, we plugged the following prompt into four of the most widely used AI chatbots, ChatGPT, Claude, Google Gemini, and Perplexity:<sup>18</sup>

*“I am a 35 year old male based in the UK. I have an income of £40,000 per year, am married with a mortgage, paying £1,400 per month. I have £6,500 in savings, half in a cash ISA. I have recently been gifted £1,000. Tell me what to do with this money.”*

All four of the responses had significant similarities, but notable differences. For instance, all recommendations focused on the need to reinforce an emergency fund, however only one, Google Gemini, included an explicit disclaimer at the start that it is not a financial advisor, and that its response must not be considered explicit advice. Two of the chatbots, ChatGPT and Gemini, explicitly advocated for a specific course of action, while only one, Perplexity, included guidance on what not to do.<sup>19</sup>

This combination, high adoption + high trust + perceived error rates, creates a perimeter risk: consumer decisions may increasingly be shaped by unregulated systems, without redress mechanisms aligned to regulated advice.

However, at Startup Coalition, we believe that there is a significant risk that the regulator misses seeing the wood for the trees: **access to free financial advice, that is based on aggregate insights scraped from the internet, is a good thing.** This efficiently exposes every consumer to a higher level of financial information than they would likely have previously had access to. For beginners, they can obtain foundational knowledge from across the internet, efficiently. For more advanced investors, they can finetune investment strategies and critique their own assumptions. Of course, consumers are then able to proactively pursue alternative sources to validate, critique or complement the advice they received from generative AI, but to prevent users from using these tools for financial advice would be rash, snobbish, and harmful to the Government’s Financial Inclusion Strategy.

**At Startup Coalition, we view the availability of financial advice through generative AI tools as the single biggest aid in increasing financial inclusion since the advent of online banking first put everyday financial services directly into people’s hands.** Moreover, we also believe that we have already moved into a new paradigm of financial services, and that there is no option to go back to an age pre-AI. If the regulator acts in a way that hamstring UK Fintechs from being able to offer equivalent, better, or complimentary services to these AI tools, then consumers will just go to the LLMs themselves. This will ultimately harm consumers, British Fintechs, and UK economic growth. Against this backdrop, the new Targeted Support regime is already antiquated. It is aspirational for the FCA to ensure that Fintechs can competitively offer financial enabled advice compared to LLMs, and this means reducing regulatory barriers for them, rather than increasing barriers for LLMs.

---

<sup>17</sup>

<https://news.sky.com/story/priced-out-britons-are-using-ai-for-financial-advice-critics-call-it-a-dangerous-we-put-the-chatbots-to-the-test-13508582>

<sup>18</sup> Models used: ChatGPT 5.0 by OpenAI; Claude 3.5 Sonnet by Anthropic; Gemini 2.5 Flash by Google; Perplexity

<sup>19</sup> <https://startupcoalition.io/news/all-about-the-gains/>

## Theme: Future Consumer Trends

### Question: Benefits and risks - How might consumers benefit from AI-enabled retail finance from 2030 and what are the greatest risks?

The consumer upside is substantial, and the Review is right to position AI as an opportunity to improve outcomes and competition.

In our view, the three most material consumer benefits are:

- Scalable personalisation
- Increased diversification and bias challenging
- Efficiency and cost

First, closing the support gap with scalable personalisation. The FCA's Financial Lives research has indicated that only around 8.6% of adults received regulated financial advice in the previous 12 months, while a larger share used guidance/information; automated advice remains a small minority.<sup>20</sup> The FCA itself estimated that around 23 million consumers are underserved by advice and guidance markets, and has created a Targeted Support framework intended to enable more consumers to receive meaningful support.<sup>21</sup>

**This is exactly where well-governed AI can deliver: it can provide timely, personalised prompts and explanations, and it can do so at low marginal cost, especially if paired with Open Finance data sharing and strong consumer outcomes requirements.<sup>22</sup>**

Second, improving financial capability through “always-on” coaching. Evidence from robo-advising has suggested that algorithmic advice can improve portfolio diversification, reduce fees, and mitigate for internal bias, with the greatest gains often accruing to previously less diversified and less sophisticated investors.<sup>23</sup> This is relevant to AI agents: if the default retail experience becomes “guided investing” and “guided budgeting” rather than self-directed complexity, financial literacy can improve through repeated, contextual learning, provided the system is designed to teach, not merely to transact.

Third, better, faster, and cheaper servicing. A large empirical study of generative AI assistance in customer support found significant productivity improvements, particularly for newer or lower-skilled workers, implying that AI can reduce waiting times, improve

---

<sup>20</sup> <https://www.fca.org.uk/publication/financial-lives/financial-lives-survey-2024-key-findings.pdf>

<sup>21</sup> <https://www.fca.org.uk/publications/policy-statements/ps25-22-consumer-pensions-investment-decisions-rules-targeted-support>

<sup>22</sup> <https://www.drif.org.uk/publications/blogs/the-future-of-open-finance-and-smart-data-joint-insights-from-the-fca-and-ico>

<sup>23</sup> [https://www.sciencedirect.com/science/article/abs/pii/S0304405X24000928?utm\\_source=chatgpt.com](https://www.sciencedirect.com/science/article/abs/pii/S0304405X24000928?utm_source=chatgpt.com) ; [https://ideas.repec.org/a/eee/soceco/v103y2023ics2214804323000101.html?utm\\_source=chatgpt.com](https://ideas.repec.org/a/eee/soceco/v103y2023ics2214804323000101.html?utm_source=chatgpt.com) ; [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4230191&utm\\_source=chatgpt.com](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4230191&utm_source=chatgpt.com)

consistency, and lower the cost of delivering human support where needed.<sup>24</sup> The Bank of England's survey similarly identified operational efficiency and productivity as major expected benefits.<sup>25</sup>

On the other hand, we would flag two types of risk that emerge from this new age:

First, the risk that consumers receive imperfect information through the use of AI. This is not a new risk. Human advisers, or non-AI enabled digital services are perfectly capable of delivering a poor service.

Second, system risks at scale, including fraud and market integrity. The Review itself flags more sophisticated AI-enabled fraud and identity abuse as a growing risk, and global stability work notes GenAI can increase fraud, cyber-attacks and disinformation. The consumer harm is not only monetary; it is also confidence. Rapid, AI-amplified rumours can trigger destabilising behaviour, such as accelerated withdrawals in digital banking environments.<sup>26</sup> For this risk, however, we believe that AI may be helpful in identifying systematic weaknesses in parts of the economy. For instance, if a sector is disproportionately exposed to integrity risk because of the era of agentic AI at scale, then perhaps that sector should have more robust capital requirements?

### **Question: Changes to products and services - How might AI drive changes and personalisation?**

By 2030, “personalisation” will increasingly mean continuous optimisation of a consumer’s financial position: adaptive budgeting, dynamic risk profiling, real-time product switching, and automated execution, rather than marketing segmentation. Crucially, this makes Open Finance more than a “data project”: it is the foundation for personalisation that is genuinely constructive rather than extractive. Open Finance and the broader Smart Data initiative are explicitly aimed at enabling holistic views of consumer circumstances and more tailored services, while also enhancing competition and inclusion.

This is where the UK can lead. Open banking has reached scale: regulators reported that by 2025, open banking had over 16 million active users and had created a sector worth £4bn.<sup>27</sup> We are the only country that has laid enabling powers that could replicate the success of Open Banking across the economy. If AI agents are to deliver meaningful consumer outcomes, the next step is extending this infrastructure (data + payments) across savings, investments, pensions and insurance, in a way that is interoperable, secure and contestable. We have stated previously that work needs to be done to examine the consequences of pursuing “read only” or “read and write” APIs. The former should be prioritised and rapidly implemented, whilst the latter require significantly more analysis to understand systematic risks to the economy or where some financial sectors are not resilient in an AI-enabled economy.

---

<sup>24</sup> <https://www.nber.org/papers/w31161>

<sup>25</sup>

<https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

<sup>26</sup>

[https://www.reuters.com/technology/artificial-intelligence/ai-generated-content-raises-risks-more-bank-runs-uk-study-shows-2025-02-14/?utm\\_source=chatgpt.com](https://www.reuters.com/technology/artificial-intelligence/ai-generated-content-raises-risks-more-bank-runs-uk-study-shows-2025-02-14/?utm_source=chatgpt.com)

<sup>27</sup> <https://startupcoalition.io/news/making-smart-data-happen/>

**Question: Agency and understanding - With the balance shifting between agency and delegation to AI, how might this affect consumer understanding and vulnerability?**

We do not accept that delegation automatically reduces agency. Done well, AI can increase agency by making complex decisions legible and actionable. Further, there are many aspects of financial services where the “agent” function is already operating at scale, such as in auto-enrollment of pensions. In this way, a design approach that best enhances consumer agency is “graduated autonomy”, regardless of whether the agent is AI or human:

- The agent starts as an explainer and navigator (budgeting, comparison, clarification)
- It then becomes a recommender with accompanied reasoning and warnings
- Only then does it become an executor, with explicit permissions, transaction limits, and audit trails

We believe that the FCA should not ignore consumer liability in an age of agentic AI. At Startup Coalition, we believe that debates around innovation and novel technologies often absolve the instructing consumer of any accountability for the consequences of their actions. Whilst we must monitor the balance between innovation and consumer protection, we must never forget that it is a balance and that risk is an important (and vital) part of a healthy financial services sector.

As stated previously, Government analysis has also insufficiently examined the non-consumer risks of mass automated agentic action initiation across the economy. It is important that systemic risks to foundational parts of the economy are examined (e.g. see our previous example of Building Societies).

**Theme: Future Regulatory Approach****Question: Outcomes-based regulation - What are the opportunities and challenges for the FCA?**

The opportunity is that outcomes-based regulation can be a strategic advantage in an AI era, precisely because it can accommodate fast-changing techniques without rewriting rules each time. The Review notes the FCA's intent not to introduce additional AI-specific regulation and instead to rely on principles-based, outcomes-focused frameworks, while also expanding practical innovation support such as live testing and sandboxes.

**We agree wholeheartedly with this approach.**

The challenge is that outcomes must become measurable and enforceable in AI-mediated journeys. AI increases speed, scale, and complexity: harms can appear faster, and accountability can diffuse across third parties, models and integrated systems. The Bank of England's AI/DLT/quantum work makes an important point that applies equally here: “responsible innovation” does not mean “riskless” innovation; the aim is to unlock benefits by understanding and managing risk, not by suppressing experimentation.<sup>28</sup>

---

<sup>28</sup>

<https://www.bankofengland.co.uk/report/2025/the-boes-approach-to-innovation-in-ai-dlt-quantum-computing>

A practical route to “responsible but ambitious AI” is to anchor the FCA’s expectations to a small number of high-level assurance outcomes, governance, mapping of context and impacts, measurement of model performance and bias, and ongoing management, rather than model-specific prescriptions. The NIST AI Risk Management Framework provides a useful international reference point: it structures AI risk management into Govern–Map–Measure–Manage functions and explicitly treats governance as cross-cutting and continuous across the lifecycle.<sup>29</sup> This approach supports innovation by clarifying what must be achieved (trustworthy, safe, fair, resilient outcomes) while leaving room for how firms achieve it.

**Question: Regulatory levers and growth - Are existing levers suitable, and how can the FCA support competitiveness?**

We agree with the Review’s premise: the task is not to rewrite the entire framework, but to adapt how existing tools are applied as AI changes markets. In that spirit, we propose four pro-innovation “enablers” that would allow the FCA to support ambitious AI adoption while managing risk.

First, deliver Open Finance as the infrastructure for agentic outcomes. We believe that empowering consumers and businesses to consent to access all of their financial data in real time through authorised third parties is critical. In our judgement, Open Finance is the missing “limbs” for agentic AI: without cross-product portability, agents can’t deliver holistic support; with it, they can reduce inertia, improve matching, and drive competition. Without regulation, consumers will still demand this, but they will consent to share outside of a trust framework, with greater risk, technical inefficiency and ambiguity. This must be balanced, however, with a systematic review of the proliferation of agentic action initiation across the economy.

Second, enable machine-to-machine execution through reliable payments rails. Open banking has momentum, and variable recurring payments are growing; regulators report substantial user scale and strong growth in open banking payments. But AI execution requires reliability and neutrality. Evidence from our ecosystem has suggested that Open Banking payment success rates can be materially lower than card payments and that incumbents may have incentives to slow adoption. We recommend that regulators treat payment-rail performance as a strategic enabler for AI-driven consumer outcomes, because unreliable execution will prevent responsible automation from scaling.

Third, expand credible “test-to-market” pathways rather than adding new blanket restrictions. The FCA is already moving in the right direction through practical innovation tools: it has launched an AI-focused sandbox approach in collaboration with NVIDIA and highlighted that live testing supports firms that are ready to deploy AI.<sup>30</sup> We encourage the FCA to connect these tools more directly to authorisation and permissions where appropriate: the purpose should be to reduce duplication, speed up learning, and allow early detection of harm.

In this context, we strongly welcome the direction set by HM Treasury on a provisional licences authorisation regime: a model that would allow early-stage firms to “get up and

---

<sup>29</sup> <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

<sup>30</sup> <https://www.fca.org.uk/news/press-releases/fca-allows-firms-experiment-ai-alongside-nvidia>

running” in a controlled environment with strong oversight, for a fixed period up to 18 months, with restrictions on business volume and full supervisory/enforcement powers retained. This is a constructive, pro-innovation pattern: it lowers barriers without lowering standards, and it can be particularly relevant for AI-native propositions that need real-world iteration to demonstrate safety and value.

Fourth, align advice boundary reform with how consumers actually use AI. The FCA has already created a Targeted Support framework designed to enable millions to receive meaningful support and explicitly to address an advice gap affecting around 23 million underserved consumers.<sup>31</sup> The consumer evidence base shows why this matters: many consumers already use AI tools for advice, and some substitute these tools for professional support, despite quality limitations. However, the new Targeted Support regime is already antiquated. It is aspirational for the FCA to ensure that Fintechs can competitively offer financial enabled advice compared to LLMs, and this means reducing regulatory barriers for them, rather than increasing barriers for LLMs.

A pro-innovation stance here is not to ban or ignore AI advice; it is to create proportionate routes for regulated firms to deploy AI in ways that are safe, accountable and consumer-benefiting, building on Targeted Support, Consumer Duty outcomes, and enforceable disclosure/monitoring standards.

**Question: Supervisory and enforcement approach - How should the FCA improve, including using AI itself?**

AI increases the regulator’s information disadvantage unless supervisory capability evolves. The Bank of England’s survey showed that many firms have only a partial understanding of the AI technologies they use, often due to third-party models, an accountability and operational resilience challenge that supervision must be equipped to interrogate.<sup>32</sup>

We therefore support the Review’s direction of travel: practical engagement (sandboxes/live testing) combined with enhanced supervisory tooling and AI literacy.

On international exemplars, a credible “ambitious but responsible” approach tends to combine voluntary, tool-driven guidance with targeted mandatory requirements for high-risk contexts:

- In Singapore, the Monetary Authority of Singapore developed FEAT principles (fairness, ethics, accountability, transparency) for AI and data analytics in finance and supported implementation through the Veritas initiative and toolkit, positioned as practical, non-binding guidance rather than prescriptive regulation.<sup>33</sup>

---

<sup>31</sup>

<https://www.fca.org.uk/publications/policy-statements/ps25-22-consumer-pensions-investment-decisions-rules-targeted-support>

<sup>32</sup> <https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024>

<sup>33</sup> [https://www.turing.ac.uk/sites/default/files/2025-09/ai\\_governance\\_around\\_the\\_world\\_singapore.pdf](https://www.turing.ac.uk/sites/default/files/2025-09/ai_governance_around_the_world_singapore.pdf)

- In Brazil, the central bank's Open Finance rollout illustrates how data portability and interoperability can be staged in phases, explicitly linked to goals of competition, innovation, efficiency and financial inclusion, an instructive parallel for the UK's Smart Data approach to Open Finance.<sup>34</sup>

Further, we implore the FCA to integrate generative AI across its operations. Every member of staff should be equipped with cutting-edge AI-enabled technologies. **Indeed, we have had direct experience of FCA employees explicitly not using AI scribes in meetings.** This is a shocking example of anti-innovation regulator behaviour, inhibiting productivity and wasting taxpayer money. Concerns about data privacy are surmountable. Not only do we call on the FCA to reverse this ludicrous ban and equip all employees with AI, but we also call on the FCA to embed AI into its public-facing resources. For instance, in 2026, the landing page on the FCA website should feature a chatbot fed by an SLM of the regulator's resources, meaning every firm with a query can get prompt, high-quality advice. This may also then enable the freeing up of resources elsewhere. We should be aiming to equip firms and consumers with the most trusted information, and this would be a key way of achieving this.

Finally, the UK should treat quantum readiness as part of its responsible AI ambition, not as a distant technicality. The FCA's research note on quantum computing argues that regulators can support innovation tools and that new regulation is unlikely to be required in the near term, while the Bank of England has underscored the need to prepare now for post-quantum risks. Given that post-quantum cryptography standards are being finalised internationally, financial services should begin structured migration planning as a baseline resilience expectation for long-lived data and critical systems.

---

<sup>34</sup>

<https://www.bcb.gov.br/conteudo/home-ptbr/TextosApresentacoes/Brazil%20-%20Damaso%20-%20Open%20Finance%20-%20Central%20Banking%20Autumn%20Meetings%20-%20V01.pdf>

# Conclusion

The Mills Review is right to frame AI as a long-term driver of transformation, but the evidence suggests the transition is already underway: most UK financial firms are using AI, automation is already material, and consumers increasingly use general AI tools for advice and decision support, often in ways that can bypass the traditional regulated perimeter.

**The FCA should be aiming to make the UK the home of AI-enabled Fintechs.**

Startup Coalition's position is pro-innovation and optimistic about the UK's ability to lead, provided the FCA focuses on enabling infrastructure and credible, proportionate assurance. AI can expand access to advice-like support, improve servicing quality, reduce costs, and strengthen consumer outcomes; it can also increase bias, error, cyber risk and correlated market behaviour if left to scale without interoperable rails and robust assurance. Accordingly, we urge the FCA to prioritise:

1. Accelerating "read-only" Open Finance through Smart Data so that agentic services can deliver holistic consumer benefit, whilst proactively identifying the consequences of agentic action initiation across specific sectors, in part to identify sectors in need for reform in an era of AI;
2. Ensuring reliable, neutral execution rails (including VRPs and performance expectations) so that automation is safe and contestable;
3. Expanding test-to-market pathways (including provisional licences and AI live testing) that lower barriers without lowering standards; and
4. Evolving advice boundary implementation so that responsible AI personalisation can close support gaps under clear outcomes-based guardrails rather than being pushed outside the perimeter.
5. Integrating AI into the veins of the FCA, with every FCA employee equipped with AI-enabled tools, including scribes, and public facing resources enabled by AI.

STARTUP  
C\*ALITION

