

The Financial Conduct Authority (FCA)
Head Office
12 Endeavour Square
London
E20 1JN
United Kingdom

Consultation Response

Which? response to the Financial Conduct Authority's (FCA) Review into the long-term impact of AI on retail financial services (The Mills Review)

Submission date: 03/03/2026

Summary

Which? welcomes the opportunity to contribute to the Financial Conduct Authority's (FCA) [Review into the long-term impact of AI on retail financial services \(The Mills Review\)](#), and ongoing engagement. As AI applications in financial services evolve at pace, we are committed to ensuring the consumer voice is at the heart of the use of AI in retail financial markets. We are exploring more ways to meaningfully engage consumers on new AI-driven financial services to understand their views.

Our response highlights the following core issues:

- **Consumer benefits and risks:** We recognise the potential consumer benefits from the use of AI tools in retail financial services if those tools are designed for the benefit of consumers, such as reducing search costs, and improved product-matching. However, we also recognise the risks to consumers, including AI tools that generate poor-quality information, misleading information or enable exploitative personalisation.
- **Consumer protection:** We support the FCA's use of sandboxes to foster AI innovation, provided they serve a dual purpose: enabling development while rigorously identifying regulatory blind spots. Sandboxes must be used to stress-test existing frameworks against AI applications including autonomous AI agents. For instance, the 'consumer standard of caution' may become obsolete if a consumer is no longer the primary actor, and some vital protections of the Consumer Credit Act could be undermined if AI intermediaries disrupt the traditional debtor–creditor–supplier chain.
- **Exclusion and vulnerability:** In addition to considering traditional financial and digital exclusion, the FCA must be mindful of new forms of disadvantage caused by data gaps of different kinds (including for consumers with high privacy standards who

opt-out of data-sharing). The FCA should also expand the risk factors of consumer vulnerability to consider low levels of AI literacy.

- **Regulatory join-up:** The development and deployment of AI in retail financial services cannot be governed in isolation. Effective coordination is imperative between the FCA and regulators such as the ICO on data privacy issues, and regulators such as the CMA on consumer protection and digital competition.
- **Regulatory approach:** As new and innovative AI products may quickly lead to consumer harm, the FCA should carefully consider whether it has the powers and toolbox to intervene promptly to address those harms.

Full response

Theme 1: future evolution of AI technology

Q2: 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)?

Focusing specifically on accountability, the shift toward agentic AI makes it easy to imagine scenarios in which protection gaps arise. For instance, well-established consumer protections, such as Section 75 of the Consumer Credit Act (CCA), may no longer function as intended. In particular, if a consumer uses a third-party AI personal shopping agent, a bank may argue that the debtor–creditor–supplier chain is broken. Similar to certain transactions involving unconnected payment processors (such as some PayPal payments), if the AI agent acts as a middleman between the consumer and the supplier, the credit provider may claim that its obligations under Section 75 no longer apply.

The FCA should proactively map possible new business models against existing protections to ensure that consumer rights are not 'designed out' by autonomous intermediaries.

Theme 3: future consumer trends

Q1: How might consumers benefit from AI-enabled retail finance from 2030 and what do you foresee as the greatest risks for consumers?

AI tools in retail finance may be beneficial for consumers. Primarily, we believe AI may help consumers to buy products that offer better value and/or are better suited to their needs if those models have been designed to benefit consumers. AI could reduce consumer search costs so that the effort required to find the best financial product is lower. AI agents may be able to instantly search and compare across providers to find the most favourable mortgage rate, savings account, or insurance premium, potentially tracking rates over time to execute a switch at the optimal moment.

Further, it could help consumers to make better choices either because AI assistants help consumers overcome behavioural biases, such as overcome inertia and avoiding impulse or fear-based decisions, or AI could be used to elicit deeper, more nuanced consumer financial preferences (e.g. risk tolerance, specific retirement goals) to match consumers with complex, appropriate products that are truly 'better' for their individual circumstances. Taken together, these factors would also incentivise businesses to innovate as there would be increased demand for new products.

However, we have identified a number of risks that could lead to significant consumer detriment depending on how AI tools are developed and implemented, and how effectively consumer protection rules are designed and enforced. We explore these risks more fully in answers to subsequent questions, but here provide a brief overview.

The first set of risks relate to the quality of information, guidance and advice provided by an AI tool. If an AI agent's objectives are misaligned with a consumer's then instead of pursuing the best financial outcome for the user, the agent might promote high-cost products to maximise commission or commercially-aligned products to maximise revenues in vertically integrated providers. Alternatively, if AI tools are of inadequate quality then they may recommend fraudulent products, while if LLMs remain prone to hallucinations then inaccuracies could lead to significant monetary detriment. Whatever the cause of bad information, guidance or advice, many consumers may be uncomfortable challenging an AI agent's recommendations/decisions, especially given the complexity of financial products, so appropriate protections are needed.

Next, the opportunity for personalisation, that could be a potential benefit, might result in greater inequality in outcomes and some consumers having worse outcomes than now. In theory, AI agents with access to vast personal data could enable highly personalized pricing or product offerings. At the extreme, this could lead to financial exclusion, but it may also cause financial harm as individual consumers are offered specific loan rates or insurance premiums that are less competitive than those offered to others. Whilst financial products can be personalised already, the extent of price personalisation is [already controversial in insurance markets](#). 'Hyper-personalisation' may also make it even more difficult to identify and challenge systemic unfairness as it will be impossible to detect whether individuals are suffering unfair bias when prices for different individuals cannot be compared.

Finally, there are concerns that the existing regulatory framework will need to be modernised as consumer journeys evolve in AI-enabled retail finance. For example, it seems unclear whether it will be possible to access redress for poor decisions involving AI as existing laws focus on protecting a human 'consumer' and may not apply to transactions executed by an AI agent. In this context, one of the challenges will be around the allocation of liability across the value chain of users, AI agents, AI developers, financial service providers and other actors.

We welcome the FCA's use of tools such as sandboxes to support innovation with AI. We believe these should also be used to identify regulatory gaps where new protections may be required to ensure innovations benefit UK consumers and do not impose unmanageable risks on them. Consumer outcomes must also be monitored, across both the intended

consumer benefits and any possible consumer harms. Our [Consumer Insight tracker](#) regularly identifies consumer data security and privacy, and business use of AI, as areas where consumers' concerns score particularly high.

Q2: Which consumer segments might 'win' or 'lose' in this new world of AI-enabled retail finance?

It will be important that the FCA considers risk factors for both traditional financial exclusion, traditional digital exclusion, and also any compounding risks around their intersection.

In addition, the FCA must be mindful of new forms of exclusion that do not follow traditional patterns. AI-enabled retail finance may replicate market disadvantages experienced by consumers about whom there are data gaps because of being new participants in a market or being under-served in that market. But AI-enabled retail finance may also create new disadvantages for consumers who are established in a market but who choose to maintain higher standards of privacy. In [our response](#) to the ICO's proposals for weakening enforcement of the regulation of online tracking (PECR), we noted that in a data protection regime where unconsented online tracking is normalised unless there is risk of discrimination, people who reject online tracking might be inferred to be in a vulnerable position, and so might experience discrimination anyway.

Q3: How might AI drive changes and personalisation in products and services, and what impact will evolving consumer expectations have? This could be to do with evolving price, value, fraud, security, mis-selling, advice, or other topics pertinent to you.

New products and services built around emerging AI capabilities may drive consumer expectations in the market in three ways that the FCA should be mindful of: consumer evaluations of quality; consumer understanding of their options when exploring a new market; and consumer experiences of 'positive friction' to support better consumer outcomes.

In the legal services sector, research has found that [lay people prefer advice generated by large language models \(LLMs\)](#) over advice generated by lawyers, because LLM syntax is typically shorter and more simply expressed than communication by legal professionals. However, advice generated by LLMs on legal issues has [consistently been found](#) to have significant errors, including 'hallucinations', which are difficult for a layperson to detect because of the fluent use of language. This may signal that in highly technical fields such as financial services, consumer evaluations of quality may not always be accurate.

Furthermore, Which?'s research on AI tools being used in consumer queries [included financial services queries](#) which yielded erroneous advice on ISA allowances and advice on HMRC tax refunds that prioritised premium paid-for services over HMRC's free tools. This raises the possibility that if consumers do not already have a sense of the market or the range of options available to them at the start of their exploration, they may not be well-positioned to critically assess the outputs of an AI tool.

Finally, while AI tools in retail financial services might yield significant benefits around decreasing friction and increasing efficiency in transactions, care must be taken to maintain norms of '[positive friction](#)' - that is, pauses in transaction processes that support better outcomes for consumers by introducing more checks for security and privacy, and more opportunities for the consumer to re-assess a transaction before completing it.

Q4: With the balance shifting between consumer agency and delegation to AI, how might this affect consumer understanding, financial literacy and vulnerability?

We welcome the FCA's consideration of behaviours and attitudes such as digital confidence, financial resilience, scam exposure, and delegation style, and how these might affect consumer agency and consumer delegation to AI.

We believe it is also important to consider how AI literacy might interact with consumer understanding and with financial literacy. [Lower levels of financial literacy](#) in the UK are attributed to socialization and lack of experience with formal financial institutions. AI tools may help overcome these specific social barriers in financial literacy by providing an alternative form of socialization in financial tools. But they may also introduce new consumer vulnerabilities or new risks around consumer understanding if consumers don't know how to use AI tools effectively. Initial analysis from DSIT indicates that [on average only 17% of UK adults](#) can explain what AI is in detail; and that marginalised groups face [greater challenges in building AI skills](#).

Q5: 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?

There are two major risks related to consumers delegating decisions over their finances to AI: AI products recommending fraudulent products, and AI products purchasing/investing in fraudulent products autonomously.

In both cases, mitigation of the risk depends on two main factors: the ability of an AI agent to distinguish fraudulent products from genuine ones, and the ability of service providers to keep fraudulent vendors off their platforms. The former mitigation will require AI financial services agents to be trained to recognise the signs of potentially fraudulent products and instructed not to recommend and/or invest in such products. The latter mitigation will rely on service providers to conduct robust Know Your Customer (KYC) checks to ensure that fraudsters are not able to access their platforms as possible vendors. KYC checks might be rendered less effective by the use of AI to create synthetic identities, which can be used to bypass such checks; service providers will likely need to factor this into the KYC checks they are performing.

On AI financial services agents potentially recommending fraudulent products: financial advisors owe their clients a duty of care in contract, under common law, but also the FCA Principles for Business (made under powers given to the FCA by the Financial Services and Markets Act 2000) include provisions for firms to perform their work with reasonable care

and skill. The Consumer Rights Act 2015 also requires every contract to supply a service to be treated as including a term that the trader must perform the service with reasonable care and skill. If financial advisors fail to do this and consumers suffer a loss as a result, consumers may be entitled to compensation. Under the current legal framework in the UK, AI agents do not have legal personhood and legal responsibility remains with the human or corporate actors who deploy and control the AI system (for further commentary, see [here](#)). There is a risk that consumers do not clearly understand the liability regime if it is not clearly communicated to them upfront, meaning that consumers may not be able to access compensation that might be owed to them should an AI agent invest consumers' money in a fraudulent product or service. To counteract this, financial services firms offering AI agents should make it clear that they are liable for any fraudulent transactions conducted by the agent and should make it straightforward for consumers to raise concerns about any potentially fraudulent transactions which the agent has conducted.

An AI agent which invests a consumer's money fraudulently will likely cause a large amount of harm to that consumer. Investment fraud is associated with especially high levels of financial harm: consumers lost £144.4 million to investment scams in 2024 and were only reimbursed in 50% of cases, well under the overall reimbursement rate for authorised push payment fraud of 88% (all data courtesy of [UK Finance](#)). The low reimbursement rate can be explained by the fact that investment scams often fall outside the remit of the Payment Services Regulator's mandatory reimbursement scheme, due to the fact that victims' are often induced to transfer funds to foreign bank accounts or into crypto wallets controlled by the fraudster. The low reimbursement rate also means that investment scams are associated with a high degree of time harm; consumers who wish to be reimbursed will often have to make a complaint to the Financial Ombudsman Service, a process which [can take months to resolve](#). [Which? research](#) has also shown that being a scam victim of any kind is associated with significantly lower levels of life satisfaction, lower levels of happiness and higher levels of anxiety. Finally, the [Payments Systems Regulator](#) has demonstrated that scam victimisation can have a negative impact on consumers' trust, including their willingness to try new approaches to managing their finances. None of these harms are specific to AI agents, but use of AI agents may contribute to these harms moving forward.

AI agents in financial services could also undermine the existing protections consumers enjoy under the Payment Services Regulator's [mandatory reimbursement scheme](#). This is because even in cases where victims are eligible for reimbursement, it is not clear currently how the use of AI financial services agents would interact with the requirement that consumers exhibit the 'consumer standard of caution' under the scheme. Should the consumers' bank interpret the use of certain AI-driven financial services agents as grossly negligent, consumers who lose money to such scams might not be covered and consequently fail to access reimbursement, creating financial harm and weakening protections in this area. The PSR/FCA should issue guidance to clarify how the use of AI agents interacts with the consumer standard of caution so that consumers who are eligible for reimbursement and not denied it on the basis that they used such an agent.

There is also a broader risk around the widespread delegation of investment decisions to AI agents: fraudsters will attempt to impersonate legitimate financial services providers. This is already a known tactic, see for example [Which?'s reporting on the 'Quantum AI' scam](#). However, the fact that AI agents are relatively new might lead consumers to be reticent

about engaging with a purported AI-driven financial services tool. Should AI-powered financial services products become more mainstream, consumer caution around such products may decrease, which could result in more consumers being defrauded by scammers pretending to be AI-powered financial services firms.

Q6: What might help make AI-driven decisions more understandable and trusted by customers, including how the use of AI may be monetised?

Current minimal requirements for robust and transparent use of AI by business is provided by data protection regulation (governed by the ICO) and general consumer protection regulation (governed by the CMA and the FCA in their respective areas).

Data protection regulations state that:

- under [the principle of transparency](#), appropriate and timely privacy information is provided to people to ensure they are sufficiently informed how their personal information is processed by AI systems;
- under [the principle of statistical accuracy](#), appropriate and effective measures in place to ensure that where appropriate datasets are relied on for determining statistical accuracy, they are accurately and fairly labelled, and do not cause detriment to people;
- under [the principle of fairness](#), fairness is a primary consideration throughout the design, development, and deployment of AI systems or components and associated personal information processing;
- and [under the principle of human review](#), AI decisions involve meaningful human review and checks, where appropriate, to mitigate eroding of privacy through selection bias and attempts to spoof controls or circumvent privacy measures.

In theory these data protection regulations should support consumer trust and understanding in AI-driven decision-making. However, the ICO's [most recent quarterly performance scorecard](#) repeats a trend of low confidence from data subjects, and long delays in addressing data protection complaints. The impact of these delays risk becoming even more significant in AI use-cases where the speed of data flows within AI systems, and the way an AI use-case can create new data about a person, may disseminate or replicate a privacy harm throughout the system or systems in which the AI model is deployed.

Consumer trust in AI decision-making also depends on robust and reliable consumer protections and effective markets. Current consumer protections prohibit '[misleading omissions](#)' that might otherwise change a consumer's decision about a product or service. In this context, business transparency about how a firm is monetising consumer interaction with an AI tool may be information that is relevant to consumer decision-making about whether to use the tool. But transparency is only meaningful if the consumer can find credible alternatives in the broader market. Without a breadth of business models and technology options, transparency about business monetisation of AI will not be sufficient for enabling consumers to act on that transparency. [Our research shows](#) that consumers often feel powerless to challenge or negotiate with organisations that hold and use their data, constrained both by limited knowledge of data practices and by their dependence on

data-driven products. These challenges are compounded by competition issues, pushing consumers into a space of 'rational disengagement' with digital ecosystems.

Finally, it is not yet clear whether current consumer protection will be adequate in agentic AI use-cases.

Theme 4: future regulatory approach

Q1: What are the opportunities and challenges for the FCA in ensuring an outcomes-based approach to retail regulation in an AI-enabled FS industry?

Q2: Are the key FS regulatory levers (Consumer Duty, Operational Resilience, SM&CR, Critical Third Party regime etc) suitable to manage future risks and to enable firms to fully take advantage of AI?

In the context of AI, it is easy to see that a legislative or rules-based regime could become obsolete if the pace of innovation exceeds the ability of legislation to keep up or stay relevant. Therefore, in terms of consumer protection, the outcomes-based approach of the Consumer Duty appears more sensible. It allows for greater technological neutrality, as it focuses on whether the consumer received a good outcome rather than whether they were served in a specific way. This approach could also help the FCA fulfil its principle of proportionality, ensuring that regulatory burdens and restrictions are appropriate to the level of harm a firm could bring to the market.

However, given the fast-paced development of AI and the innovation that is expected to follow, it is possible that a lack of clear prescription in rules or legislation could lead to confusion. For example recently, there has been industry criticism of [decisions made by the Financial Ombudsman in favour of consumers](#). Firms have argued that they [lacked clarity](#) about the intentions of the FCA's regulations and about which behaviours were or were not intended under the FCA's outcomes-based approach.

To enable firms to take full advantage of AI, greater clarity through more frequent updates to relevant guidance would likely be helpful. Generally, the FCA is required to consult on changes to guidance, as well as to the Handbook rules. However, there are exceptions to the obligation to consult. For example Section 138L of FSMA provides a specific 'urgent changes' exemption that allows the FCA to bypass these requirements under certain conditions.

Given that new and innovative products may quickly lead to consumer harm, the FCA should carefully consider whether it has the power to provide clarity to firms promptly, for example through urgent guidance, when harm has arisen. It should also consider whether this mechanism, or others, would allow it to issue guidance repeatedly in this way without undermining its legitimacy in the longer term.

Q3: 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?

AI is already in use in order to assist front line customer service agents across a number of sectors. This includes co-pilot tools that can conduct sentiment analysis which can alert an agent in real time if a consumer is becoming frustrated. We can expect its use and proficiency to develop substantially by 2030.

Given the apparent [capability of AI to detect some markers of vulnerability](#) that may not be easily detectable to people, the FCA should consider how it expects firms to adopt AI to assist in identifying consumers with indicators of vulnerability in both the immediate and the long term.

While capabilities in language processing and sentiment analysis have improved, their effectiveness is still unproven on a mass scale and technical biases could currently lead to poor consumer outcomes. While the potential for AI to identify drivers of vulnerability at scale offers an opportunity for firms, the FCA must remain vigilant against certain risks such as tools that lack a robust scientific basis, context blindness or cultural bias, and the possible erosion of human accountability where employees may stop exercising sufficient professional judgement.

That being said, over the longer term as these technologies develop the FCA should consider amending its expectations to make clear clarify that relying solely on employee observation or customer self-disclosure may no longer be sufficient to meet the Consumer Duty's 'foreseeable harm' standard, particularly where firms using such methods consistently fail to identify vulnerability as required. Instead, firms should be encouraged to consider technology-driven behavioral and sentiment analytics as an augmentation of its consumer facing services.

If these tools prove effective this could help reduce the 'disclosure gap'. At present the identification of vulnerability often relies on the customer's willingness or ability to self-identify. However, many consumers (particularly those experiencing mental health crises or cognitive decline) do not recognise their own vulnerability or may feel embarrassed to disclose it.

Some AI tools are already in development seeking to better analyse speech patterns, transaction anomalies, and communication frequency to better identify drivers that a human advisor would likely miss. If the technology exists to reliably identify these risks before harm occurs, or at a point where a firm could intervene in order to reduce further harm, failing to use it could be viewed as running contrary to the expectations set out under the consumer duty.

Specifically, within the mandatory reimbursement framework for APP fraud there is an increasing focus on ensuring institutions do not overlook hidden or situational vulnerabilities when assessing the consumer standard of caution. The FCA should explore how the use of AI in the future could be expected to be used as best practice in order to assist in the identification of vulnerable consumers.

More widely, as the sophistication of AI-driven fraud increases, the FCA (once the responsibilities of the PSR have been absorbed) should consider how this will impact its expectations regarding 'gross negligence' and the application of The consumer Standard of Caution Exception Guidance in relation to Authorised Push Payment (APP) fraud cases. The traditional standard of caution rests on an individual's ability to detect inconsistencies that are increasingly becoming challenging to observe. When generative AI can produce hyper-realistic deepfake audio and video to impersonate bank officials or family members, and legitimate payment interfaces can more easily be mimicked, the threshold for what constitutes a 'significant degree of carelessness' might need to be reformed. Failure to adjust this standard risks punishing consumers for failing to detect deceptions that could be, by design, indistinguishable from reality.

About Which?

Which? is the UK's consumer champion, here to make life simpler, fairer and safer for everyone. Our research gets to the heart of consumer issues, our advice is impartial, and our rigorous product tests lead to expert recommendations. We're the independent consumer voice that works with politicians and lawmakers, investigates, holds businesses to account and makes change happen. As an organisation we're not for profit and all for making consumers more powerful.

For more information contact:

Dr Milly Zimeta
Head of Digital & Data Policy

Paddy Greene
Head of Money Policy

c/o policyandinsight@which.co.uk

March 2026