Amplified Global Response: FCA Review into the long-term impact of AI on retail financial services

Introduction

We are pleased to respond to the FCA's Mills Review into the long-term impact of AI on retail financial services.

Our product, Amplifi, is underpinned by AI and machine learning. Amplifi has been developed using a wide range of established and proprietary linguistic models to accurately test intelligibility of written information, i.e. how likely a reader is to be able to read, understand, and use that information.

We are also developing new and innovative ways to present information in a more engaging and intelligible form using digital layering, and by testing comprehension in situ. Our existing tool uses AI to assess the text and to provide an intelligibility score and identify ways to improve it, while we also use generative AI in providing simplified text.

Our direct experience upon which this response is founded has been derived from both our development of the tool, and through its use by a range of financial services firms. They use Amplifi to assess their communication, improve and simplify them, and to report on their compliance with the FCA's Consumer Duty and the Consumer Rights Act.

Summary

Our submission particularly focuses on our direct experience, of how AI can enhance content creation, consumer understanding, and regulatory compliance in financial services.

Key Points

1. AI Adoption in Financial Services
- AI is increasingly already being used for a wide range of tasks including content creation, risk assessment, fraud detection, customer service, and regulatory compliance.
- For consumers, there is a growing need for AI tools that simplify financial information, ensuring consumers understand key documents like contracts and credit agreements. AI can also help support how they engage and find information, their level of comprehension, and increase the ability of consumers to be educated while they engage.
- For firms, there is a growing need for education (firms' internal AI users often aren't AI experts), support, governance and controls.

2. Consumer Benefits & Challenges
- Many consumers struggle with complex financial terms, leading to poor decision making.
- AI can improve financial literacy by simplifying documents, restructuring content and reducing the cognitive load for readers, and providing interactive tools for better comprehension.
- However, AI models risk bias, and human oversight is necessary to ensure fairness.

3. Barriers to AI Adoption
- Regulatory challenges, ethical concerns, and trust issues slow down AI implementation.
- AI must integrate seamlessly into financial firms' workflows and remain transparent to comply with laws and firm expectations.
- There is a growing need for sector-specific AI models, rather than relying on general AI solutions.

In addition, there are a number of practical and organisational barriers that materially affect firms' ability to adopt AI responsibly and at scale:
- Cost and economic risk remain significant constraints. Beyond upfront development costs, firms must account for ongoing expenses associated with deployment, monitoring, assurance, and model maintenance. There are also material downside risks, including the potential cost of failure, including remediation, customer redress, regulatory enforcement, and reputational damage. These issues can make firms cautious about deploying AI in regulated customer-facing contexts.
- Leadership, culture, and governance maturity are also critical barriers. Effective AI adoption requires clear senior ownership, informed decision-making, and a culture that balances innovation with risk management. In many firms, leadership teams may lack sufficient AI literacy to confidently set strategy, challenge vendors, or oversee AI-enabled processes. This can lead either to over-cautiousness or to inappropriate reliance on third-party solutions.
- Firm readiness and operational capability vary widely. Many organisations lack the internal skills, processes, and controls needed to design, test, validate, and monitor AI systems over time. This includes the shortcomings, and relative infancy of AI assurance.
- Infrastructure and data constraints further limit adoption. Legacy IT systems, fragmented data architectures, and poor-quality or insufficiently governed data can make it difficult to deploy AI effectively or safely. Reliance on external cloud and AI providers also introduces dependencies that must be carefully managed under operational resilience and third-party risk frameworks.

4. Recommendations for AI Regulation
- Regulators should focus on outcome-based approaches, ensuring AI tools meet consumer protection standards.
- Financial services should be encouraged to test AI in regulatory sandboxes before full deployment.

- Cross-regulator collaboration (e.g., FCA and legal regulators) can promote innovation while ensuring compliance.

5. Balancing AI Risks & Opportunities
- AI can enhance productivity but should not replace human decision-making in regulated content creation.
- Governance measures, such as explainability requirements and human oversight, are essential to mitigate risks.

AI has the potential to revolutionise financial services, making information clearer and improving regulatory compliance. However, ethical concerns, bias, and regulatory challenges must be addressed to ensure responsible AI adoption.

Theme 1: Future evolution of AI technology

1. AI technologies most likely to transform UK retail financial services from 2030, and why

AI technologies that we feel have the greatest potential to reshape retail finance include:

- **Advanced Generative AI\:** enabling richer natural language interaction between consumers and firms, higher-quality product and legal explanations, and dynamic regulatory compliance reporting. These can improve comprehension of complex financial information and reduce consumer asymmetries in understanding.
- **Human-centred AI interfaces**: systems designed to prioritise intelligibility and explainability. These technologies directly address barriers to consumer understanding.
- **AI assurance platforms**: integrated tools for monitoring, auditing, and validating content for compliance with consumer protection and fairness standards.

These technologies are already being deployed in different ways across the industry, including our own product Amplifi.

2. Potential and direction of agentic AI: implications for retail finance

Agentic AI (systems that can autonomously plan, reason, act, and learn to achieve complex, multi-step goals with minimal human intervention) has the potential to enhance many financial services from passive, rule-based automation to proactive, autonomous operations. Unlike generative AI, which creates content based on prompts, agentic AI can execute tasks across multiple systems.

Some of the important value that agentic AI can bring to financial services include:

1. Efficiency and Cost Reduction

Agentic AI has the potential to automate complex, end-to-end workflows that previously required significant human expertise, resulting in productivity gains.

2. Real-Time Risk Management and Enhanced Compliance
Agentic AI can provide continuous monitoring, allowing firms to move from reactive risk management to proactive, real-time protection.

3. Personalisation and Improved Customer Experience
Agentic AI will increasingly allow firms to move beyond generic, static interactions to tailored, proactive financial coaching and services.

However, these technologies, while fast moving, are in the relatively early stages of deployment, and significant testing and adaptation will be required for some tasks and value to be realised.

There is a danger of deploying before proper due diligence and testing has been carried out.

There is also a risk of diminishing consumer and firm agency if users delegate decisions without fully understanding risks or retaining control. For instance, for consumers there is potential for opaque automated advice that feels personalised but lacks transparency or robust consumer protection.

Implications may include:
- **Liability and accountability gaps:** establishing who is responsible for harms when agentic AI acts on behalf of consumers will demand updated governance frameworks.
- **Assurance and explainability needs:** firms should embed explainability at the core of agentic system design, with audit trails ideally visible to both consumers and regulators.

3. Interaction of AI with other digital technologies through 2030

AI is likely to increasingly intersect with:
- **Open Finance and API-based ecosystems**, enabling more personalised financial products and services.
- **Digital identity protocols**, enabling richer consumer profiles but also raising privacy trade-offs.

These combinations could create new market segments (e.g. programmable finance, smart contracts managing claims adjudication), but also new controls and standards will be needed to mitigate aggregated systemic and consumer risks.

4. Impact on you: How will AI change your operating model, operating environment, and dependencies from 2030? How might you respond to wider adoption of AI?

By 2030 we anticipate:

- A significant shift toward integrated, **AI-assisted regulatory compliance solutions**, where firms routinely use AI for documentation, summarisation, and decision-explanation tasks.
- Our own model development will increasingly be audited and validated against both regulatory frameworks and consumer protection outcomes.
- Strategic alignment with AI assurance platforms to demonstrate compliance with fairness and transparency standards, a logical evolution of our current focus on intelligibility of outputs.

5. UK's comparative advantages and gaps in AI

Advantages:
- A principled-based regulatory legacy (e.g., Consumer Duty) that can adapt to AI contexts.
- Strong university research in explainable and robust AI.

Gaps:
- Limited dedicated data centre capacity relative to the US and China.
- Shortage of AI governance and assurance expertise within firms and regulators.

Recommended actions to improve competitiveness:
- Support specialised AI governance and assurance standards.
- Facilitate and encourage cross-sector sandboxes for shared learning on consumer outcomes.

Theme 2: Future impact of AI on markets and firms

1. Market structure and barriers to entry

AI could both lower and reinforce barriers to entry. Smaller firms may leverage cloud-native tools to compete, but larger incumbents with significant data and computing advantages could entrench positions through data insights and economies of scale.

Where cost savings arise, if they are to be passed on to customers in the form of reduced pricing and better value services, this will require transparent performance and outcome reporting frameworks.

2. Self-reinforcing market dynamics

Data advantages and network effects (firms with the most extensive customer interactions improving models fastest) could lead to a 'winner-takes-most' dynamic. This could in part be counterbalanced by increased standards for data portability, interoperability, and shared benchmarks, and improving sandbox data, including links to better outcome and longitudinal data.

3. Control of customer relationship to 2030

- Incumbent financial firms may maintain relationships through trusted brands.
- However, AI intermediaries and consumer agents could capture loyalty if they provide superior experience and transparency, especially among digitally savvy users.

4. Regulatory perimeter risks

Unregulated AI intermediaries could provide advice or intermediation that functionally mirrors regulated services. This may draw significant consumer interactions outside the regulatory perimeter unless clear standards for accountability, redrawing of regulatory boundaries and/or enforceable obligations are established.

Theme 3: Future consumer trends

1. Consumer benefits and risks

Benefits include enhanced personalisation, improved intelligibility and actionability of product terms, pricing and product/promotional information, and (longer term) automated optimisation of everyday financial decisions.

Risks involve opaque decision-making layers, diminished agency, and model bias potentially leading to discrimination or exclusion, particularly if consumers rely on agentic AI without human oversight mechanisms.

2. Inclusion vs. exclusion

Consumers with higher digital confidence are more likely to benefit, while those with low digital literacy or limited access to technology may be excluded. Accessibility and design must prioritise human-centred intelligibility.

3. Changes to products and services

AI will drive hyper-personalised offerings, but there must be guardrails against mis-selling and misleading outputs, particularly where AI hallucinates or over-optimises for short-term engagement.

Ensuring explainability and human-led oversight and governance of AI processes will remain an important factor for many services harnessing AI.

4. Agency and understanding

Delegation to AI agents may reduce active financial literacy unless systems are designed to improve comprehension and support informed decision-making. Tools that explain reasoning, risk vs reward and trade-offs are essential.

5. Fraud evolution

As consumers delegate more actions to AI, fraud schemes may become more sophisticated (e.g. deepfakes, synthetic identities). Industry must invest in AI-based defensive systems and adaptive fraud detection.

6. Trust and monetisation

Consumer and industry trust will rely on transparency, explainability, and clear disclosure of how AI makes decisions.

AI monetisation models should not incentivise opaque upselling or reward bias.

Theme 4: Future regulatory approach

1. Outcomes-based regulation

An outcomes-based regulatory framework is well suited to an AI-enabled financial services environment, as it allows innovation while maintaining a clear focus on consumer protection. The primary opportunity lies in validating AI systems against measurable consumer outcomes, such as improved understanding, fair treatment, and avoidance of foreseeable harm, rather than prescribing specific technologies or methods.

However, this approach also presents challenges. AI systems, particularly generative and agentic models, can be opaque, adaptive, and probabilistic in nature. This may make it harder to evidence outcomes consistently over time. To address this, outcomes-based regulation should be supported by robust testing, monitoring, and assurance mechanisms. This should enable firms to demonstrate that AI-driven processes continue to meet regulatory expectations as their models and use cases evolve.

2. Regulatory levers

Existing regulatory levers provide a strong foundation for governing AI use in financial services. Frameworks such as the Consumer Duty already emphasise customer outcomes, clarity of communications, and avoidance of foreseeable harm, which align closely with the risks and opportunities presented by AI-generated and AI-assisted content.

Similarly, SM&CR supports accountability by ensuring clear ownership of AI-enabled processes, while Operational Resilience and the Critical Third-Party regime are increasingly relevant as firms rely on external AI providers, cloud infrastructure, and shared models. These levers remain suitable, but their application will need to evolve to explicitly incorporate expectations around model governance, explainability, auditability, and human-in-the-loop oversight, particularly for generative and agentic AI systems.

3. Supervisory and enforcement approach

As AI adoption increases, supervisory approaches should place greater emphasis on governance, auditability, and accountability, rather than solely on static pre-deployment approvals. The FCA could enhance supervision by encouraging firms to maintain clear documentation of model design, training data considerations, testing results, and ongoing performance monitoring.

There is also an opportunity for the FCA to use AI tools itself to support supervision, for example by identifying patterns of consumer harm, monitoring large volumes of regulated communications, or analysing firm submissions more efficiently.

Enforcement frameworks should continue to focus on accountability, ensuring that responsibility for AI-enabled decisions and outputs remains clearly attributable to regulated entities and senior managers.

4. Growth

The FCA can support growth by maintaining a regulatory environment that is predictable, proportionate, and innovation-friendly, while remaining anchored in consumer protection. Regulatory sandboxes and live testing environments are particularly valuable for AI, allowing firms to pilot new applications, assess real-world consumer outcomes, and iterate safely before full deployment.

Clear guidance on regulatory expectations for AI assurance, explainability, and human oversight would further reduce uncertainty for firms, particularly smaller and scaling providers. By focusing on outcomes rather than prescriptive rules, the FCA can enable firms to innovate responsibly while maintaining trust in UK financial services.

5. Frameworks for inspiration

The FCA may wish to draw lessons from emerging AI assurance and governance frameworks in other regulated sectors, such as healthcare, safety-critical systems, and data protection, where transparency, auditability, and risk management are well established.

It is hard to point to any one example, as the operating environment and AI capability is rapidly changing, and the frameworks in which they operate are relatively new, and under development.

One key question for the FCA and Government policy makers is the most appropriate approach. We favour an outcome-based and evidential system, as it can remains agile while ensuring adequate consumer protection.

An effective national framework for AI is likely to combine binding, risk-based legislation and evidence-based outcome requirements, with voluntary technical standards for governance.

The EU AI Act is an example of a comprehensive legal framework. It classifies AI systems by risk (unacceptable, high, limited, minimal), and enforces strict obligations on high-risk AI.

An example of a set of voluntary technical standards is the US NIST AI Risk Management Framework. It Focuses on trustworthiness, documenting, and managing AI risks through a voluntary, flexible approach. It emphasises transparency, safety, and accountability.

An effective national framework is likely to need to combine both approaches effectively to remain nimble and future-facing, while retaining as much certainty for firms and innovators, and high levels of protection for consumers.

International coordination and inter-regulatory dialogue will also be important, particularly as AI systems and providers operate across borders.

Cross-regulator collaboration, both within the UK and internationally (for example through the GFIN Network), could help develop consistent standards for AI governance. This would reduce fragmentation, and support firms operating in multiple jurisdictions, while maintaining strong protections for consumers.

Ewan Willars
Regulatory Lead, Amplified Global Ltd
ewan@amplified.global