Institute and Faculty of Actuaries

24 February 2026

FCA: The Mills Review

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

The Institute and Faculty of Actuaries (IFoA) is a royal chartered, not-for-profit, professional body. We represent and regulate over 34,000 actuaries worldwide, and oversee their education at all stages of qualification and development throughout their careers. Actuaries are big-picture thinkers who use mathematical and risk analysis, behavioural insight and business acumen to draw insight from complexity. Our rigorous approach and expertise help the organisations, communities and governments we work with to make better-informed decisions. In an increasingly uncertain world, it allows them to act in a way that makes sense of the present and plans for the future.

Key points

- The IFoA welcomes the opportunity to respond to the FCA's Mills Review into the long-term impact of AI on retail financial services. The rise of AI offers significant potential to the financial services sector (and indeed wider society), and many of the benefits of AI are already being realised. Getting the balance right in terms of making the most of AI's upside potential whilst having appropriate mitigants to the downside is relevant to the wider public interest.
- AI is becoming increasingly embedded across the financial services sector. Actuaries are particularly involved in developments relating to general/ life/ health & care insurance and pensions. AI adoption will likely continue to accelerate due to increasing data availability, computing power, multimodality and the use of autonomous AI agents.
- We expect financial service firms will likely adopt AI-driven Regulatory Technology, enabling better compliance, automated reporting and real-time risk identification.
- AI's potential to improve productivity within the financial services sector is extensive.
- It is likely that AI will lead to considerable displacement of roles within the financial services sector. There is, however, scope for individuals to reskill and focus on alternative, value-adding roles, including interpreting AI-generated outputs, model governance, risk management and AI ethics.
- There is a challenge to reconcile the upside of AI with its extensive and growing energy use, and the resulting challenge to achieving climate goals. These potentially conflicting objectives need to be managed.
- The UK is well-placed to be a leader in responsible AI adoption, provided that regulation balances innovation with robust consumer protections, and ethical guardrails are in place. The UK's Fintech ecosystem and globally recognised AI research/ expertise should also help give the UK a comparative advantage.
- AI can increase exposure to cyber risk. AI-driven cyberattacks can exploit vulnerabilities at scale and amplify potential impacts through the interconnectedness of systems and devices. Reliance on a small number of AI providers (such as large tech firms) could potentially introduce systemic risks.
- AI offers consumers a wide range of benefits including those specific to vulnerable consumers (such as better identification of vulnerability). However, there is also a risk of AI replicating existing biases.
- We favour a broad principles-based, rather than rules-based, approach to AI regulation. We believe it is important there is adherence to global standards to ensure a level playing field.
- We believe the IFoA has an important role to play in the debate on the future evolution of AI across the financial services sector. We hope the FCA finds our response to this inquiry helpful and constructive, and we would be delighted to discuss it further.

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

Introduction

1. The Institute and Faculty of Actuaries (IFoA) welcomes the opportunity to respond to the FCA's Mills Review into the long-term impact of AI on retail financial services. As in mentioned in the introduction to the FCA's review, AI has significant potential to transform retail financial services, and this transformation is already well underway.

2. We believe the focus of the FCA's review is also helpful, in considering the potential future evolution of AI technology, its impact on the market, firms, consumers and indeed AI regulation. AI offers retail financial services a wide range of benefits, but there is also a potential downside: AI can generate or exacerbate a range of challenges. The benefits and risks are very relevant to conduct regulation. Getting the balance right in terms of making the most of AI's upside potential whilst having appropriate mitigants to the downside is relevant to the wider public interest.

3. It is important to note that, as for any IFoA response, we have considered the FCA's review from an independent, public interest perspective. In doing so we have taken a broad perspective on the public interest, considering the impact of AI and AI regulation on current and future financial services consumers, and society as a whole.

General Comments

4. AI has significant potential to promote innovation across a range of retail financial services in which our members practise, including general/ life/ health & care insurance and pensions products. However, as it delves ever deeper into our lives, AI raises significant questions over ethics and the public interest, including conduct risk concerns.

5. It could be said that actuaries are the original data scientists, and data science remains an integral and indispensable aspect of the actuarial profession today. Actuaries analyse data, form judgements and advise financial service firms. These core functions will remain key as AI evolves, with the rise of AI expanding the toolkit available to actuaries, and to regulators. Actuaries' existing skills in critical thinking, ethics, judgement and understanding model limitations will be particularly important when applied in an AI context.

6. In late 2024 we issued updated guidance for our members focused on ethical and professional use of data science and AI. The guidance recognises recent advances in AI techniques and capacity, with new and enhanced risks and opportunities resulting from this.

7. While this is not the focus of the current review, we note that one key public interest consideration is reconciling the upside of AI with its extensive and growing energy and water use, and the resulting challenge to achieving climate and resource goals. These potentially conflicting objectives need to be managed. The IFoA has done extensive work on the challenge of climate change, particularly in our recent paper on planet solvency[1].

[1] See: https://actuaries.org.uk/news-and-media-releases/news-articles/2025/jan/16-jan-25-planetary-solvency-finding-our-balance-with-nature/

8. Given the above, and the importance of collaboration, we believe that the IFoA has an important role to play in the debate on the future evolution of AI across the retail financial services sector, including in relation to conduct (and prudential) regulation. We welcome the opportunity for further engagement with the FCA to provide actuarial expertise in the evolution of retail financial services in an AI-dominated world.

9. We hope the FCA finds our response to the Mills Review helpful and constructive, and we would be delighted to discuss our response with the FCA.

10. For completeness, unless otherwise specified our response covers both narrow AI (primarily data science and machine learning) and Generative AI (GenAI), each of which will bring different risks and opportunities.

11. We have responded to the questions within the inquiry where we have specific points to raise.

FCA Review Questions

Theme 1: Future Evolution of AI Technology

Question 1: AI technologies: 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.

12. Firstly, it is helpful to consider the current reach of AI across retail financial services. AI is becoming increasingly embedded right across the financial services sector, with varying levels of coverage. Examples where actuaries are particularly involved in AI development include:

    - **general insurance**: AI models enhancing underwriting, claims processing, and fraud detection;
        - granular datasets and predictive analytics improving pricing, risk segmentation, and customer retention;
        - AI models also help with customer segmentation and target marketing campaigns in a more tailored way, to meet customer needs;
    - **health & care insurance**: AI models are being used to develop personalised disease management programmes, helping both the customer with the management of their health, and the insurer in containing their claims outgo;
    - **life Insurance**: more granular mortality and longevity predictions; underwriting triage adds efficiency;
    - **pensions**: AI can help analyse investments, assess long-term risk scenarios, assist in fraud detection and improve member engagement by providing personalised insights into retirement savings.

13. AI is being used widely across the financial services sector to assist in the preparation of financial statements, automating many manual processes and developing audit trails.

14. One downside of greater reliance on AI both in financial services (but also more generally) is potential loss of future expertise, if AI replaces humans in too many areas and the training regime for future professionals becomes disrupted.

15. Over the next five-ten years to 2030 and beyond we can envisage that:

    - AI adoption will likely continue to accelerate due to increasing data availability, computing power, multimodality (text, audio, image, video), and the use of autonomous AI agents (Agentic AI);
    - Note in particular that we expect financial service firms will likely adopt AI-driven Regulatory Technology (RegTech), enabling better compliance, automated reporting and real-time risk identification.
    - AI-driven investment tools will enhance personalised investment planning and portfolio optimisation;
    - Fintech firms will likely continue to be at the fore in AI adoption due to their agility and tech-driven models;
    - Further tech giants will likely enter the insurance/ general financial sector and compete with incumbent/ traditional players;

Question 2: Agentic AI: 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)?

16. Agentic AI is likely to lead to increased efficiencies for firms. Many firms have successfully worked with process automation for some years, but Agentic AI increases the power of this. By combining process automation with Generative AI - specifically LLMs – it permits the Generative AI tools to trigger actions as part of the whole process. This enables automation of a wider range of process steps which used to require human decision-making, particularly those involving unstructured data, such as emails and call transcripts.

17. Such automation requires careful oversight and risk management by firms. Regardless of technology used, the firm retains the accountability for the decisions taken, and particularly in retail finance must ensure transparency of decision making to be able to explain to the customer why a particular decision has been taken. Firms must therefore ensure they have adequate oversight of any Agentic AI deployment, and this should be designed in from the start, not added as an afterthought.

18. Agentic AI suffers the same risks as LLMs, including hallucination, missing out items, and bias. A further risk is prompt injection, which occurs because LLMs cannot distinguish easily between 'data' they are reading, such as an email, and new instructions, such as a hidden text string trying to influence the Agentic AI outcomes. The source of such hidden prompts may be legitimate, for example a customer trying to get the automated process to proceed in their favour. However, it could also be malign, for example a hacker aiming to exfiltrate confidential customer data, and this will be an important security concern for firms and regulators.

Question 3: Digital technologies: How do you anticipate AI combining with other digital technologies, resources, and infrastructures through 2030? What specific markets, products, controls, standards, and risks could emerge?

19. Over the period until 2030 we think the degree of interconnection between AI and other technologies may be limited, primarily because traditional financial services firms have considerable legacy technology that might not easily align with AI. Looking further ahead, it is plausible to envisage a time when Agentic AI is handling the bulk of firms' administration, and staff are not needed for traditional data processing roles. In this scenario firms are likely to need more people with responsibilities for AI oversight and risk management.

Question 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?

As we are a global membership body for actuaries, and not a financial service product provider, this question is not directly applicable to the IFoA. However, we make some observations below in relation to how some product providers could be impacted.

20. Last year the Bank of England published findings showing that 75% of financial services firms were already using AI, with a further 10% planning to use it over the next three years. It is plausible that financial service firms have adopted AI at a greater rate than other sectors, not least due to the structured data-driven nature of the industry but also due to competitive pressures. However, the existence of legacy systems may slow down the pace of adopting AI for some firms within the sector.

21. In the next paragraph we list a number of potential benefits of AI. However, such benefits are not automatic – they depend not only on using AI but on using it effectively. According to a recent PwC survey[2], 56% of companies are 'getting nothing' out of AI in terms of revenues or costs, and they

[2] https://fortune.com/2026/01/19/pwc-global-chairman-mohamed-kande-ai-nothing-basics-29th-ceo-survey-davos-world-economic-forum/

attribute this to a lack of focus on basics such as 'clean data, solid business processes, and governance'.

22. In theory, AI has significant potential to improve productivity within the retail financial services sector. Possible AI applications are quite extensive. Current and prospective examples include:

    - automating time-consuming, routine, 'low risk' tasks (e.g. timesheets, meeting notes, model documentation), thus freeing-up individuals for more value-adding/ impactful work;
    - automating repetitive tasks (e.g. claims processing, underwriting, data cleansing and assembly, compliance checks, reserving, claims analytics, experience studies);
    - AI-powered text or voice customer chatbots to attend to basic customer inquiries, perhaps using human oversight for more complex or material queries;
    - fraud and anomaly detection in financial transactions;
    - more efficient customer service by better identifying customer needs, including '24/7' customer service, when the customer wants it;
    - enhancing efficiency in the delivery of regulated financial advice, through better compliance controls and streamlined processes;
    - development of personalised user-interfaces with products and services;
    - improved understanding of product protection gaps, and identification of cheaper products, more accessible to a wider population;
    - AI-powered investment tools helping individuals better engage with their long-term savings and pension investments;
    - model coding allowing firms to replace legacy systems and large numbers of spreadsheets with more efficient, better-governed systems;
    - more refined risk assessment in insurance pricing (already quite embedded), extending to capital modelling and reserving;
    - large organisations can use AI to improve consistency across their internal reporting, and adherence to group standards;
    - AI-driven regulatory solutions automating contracts and analysing regulatory text;
    - upskilling employees by using GenAI as a tutor, including on AI content.

23. Financial service firms - in the retail sector or otherwise – will need to consider the risks associated with relying on third party AI providers, including risks relating to security and compliance. Given the application of AI tools to a firm's client/ wider data, appropriate data governance will be key.

24. Where firms buy AI tools from an external provider for use in their business, they may not have full clarity over the complexity/ data/ parameterisation et al. This potentially increases third party dependency concerns. Furthermore, not all models produce readily explainable results.

25. Using third party AI tools also runs the risk that the tool provider 'harvests' the client firm's data. Robust and clearly understood contracts will be key mitigant here.

We note that the Review scope does not include potential employment impacts of AI. However, we include the following points as we believe they are relevant to future conduct risk (and conduct risk management).

26. It is quite likely that AI will lead to considerable displacement of roles within the retail financial services sector. AI has significant potential to replace administrative and manual processing jobs. However, there could be scope for individuals to reskill and focus on alternative, value-adding roles, including in interpreting AI-generated outputs.

27. There is also significant potential for the rise of AI to create new roles in model governance, risk management and AI ethics - subject to appropriate training being provided. Data and model requirements around interpretability, explainability and transparency will require highly qualified multi-disciplinary experts with a high-level understanding of all aspects of AI (e.g. algorithms, processes, results and governance).

Question 5: The UK: 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?

28. In our view the UK is well-placed to be a leader in responsible AI adoption, provided that regulation balances innovation with robust consumer protections and ethical guardrails in place. The UK's Fintech ecosystem and globally recognised AI research/ expertise should also help give the UK a comparative advantage.

29. In the Fintech industry, many startup firms have a focus on software development. Fintech firms may therefore be better-suited to adopting AI, where they have strong technical coding expertise. Building, maintaining and improving AI tools/ models requires technically-proficient software engineering skills. Furthermore, Fintech firms are often quite agile in nature.

30. Fintech firms often operate using a cloud-based approach. This can be easier to integrate with AI systems compared with traditional financial institutions using legacy systems, that are often difficult to retire. One limitation in the UK's position is that our AI applications rely on US cloud technology and AI foundational models, and access to these could be vulnerable to political interference.

31. The development of open finance is also relevant to the UK's aspirations to be a global leader in AI. AI has significant potential to 'turbo charge' the expansion of open finance and build on the success of open banking. The need to harness technology and support innovation is particularly important given consumers' growing digital capability in recent years.

Theme 2: Future Impact of AI on Markets and Firms

Question 1: Market structure and customer passthrough: How might AI change concentration in your market? What are the current drivers of concentration, and which could AI disrupt or reinforce? Do you expect AI to increase or decrease barriers to entry? Do you expect AI to increase concentration, reduce it, or reshuffle who the dominant players are? Where there are cost reductions do you expect these to be passed on to customers as lower prices?

Market Concentration

32. The creation of suitable bespoke AI/ Large Language Models would likely be beyond even the largest of financial service firms, so there is then a need to use external providers. Fewer providers of AI tools would mean fewer algorithms on the market, which increases the risk of herding behaviour. This could be mitigated by regulatory approval and monitoring of models used in financial services, and potentially intervention to help ensure there is a wider choice of models. However, such intervention needs to reflect the global nature of AI and the environment in which large tech firms operate.

33. Currently there are (say) five external AI model providers. All except Chinese model DeepSeek are US-based. Nevertheless, this exposure is less concentrated than the current reliance on Microsoft for Excel operating systems and base programmes.

34. There is an emerging trend for larger insurers to move away from publicly available AI (such as ChatGPT) to more bespoke solutions, trained on their own data. For now, bespoke AI solutions are less affordable for smaller insurers, and in any case, they are still based on the original models rather than being new own-firm LLMs.

35. Reliance on a small number of AI providers (such as large tech firms) can introduce systemic risks. Mitigations to such systemic risk could include performing independent AI audits and regular reviews to prevent hidden biases and systemic risks. AI models could also be subjected to stress tests and explainability requirements. Systemic risks could include wider market disruption from failure of significant systems, as well as potential security issues which could be exacerbated when geopolitical tensions arise.

36. In theme 1, question 2 we mentioned that Agentic AI can enable greater automation of processes which used to require human decision-making, such as those involving unstructured data such as emails and call transcripts. This will generate cost savings for firms, and we expect those firms would want to pass those savings back to customers in order to take a larger share of the market.

Barriers

37. There is a range of potential barriers impeding the adoption of AI in financial services. Key barriers include:

    - **cost** - AI transformation projects can often be expensive, with costs arising in relation to consultancy advice implementation, IT infrastructure and staff training;
    - **workforce training and adaptation** - AI integration necessitates upskilling financial professionals; which requires investment in terms of time as well as cost;
    - **data quality and availability** - AI models require high-quality, unbiased data if to be used with confidence;
    - **regulatory uncertainty** - firms may be cautious to embark on AI development due to unclear AI governance frameworks;

    - concern that regulators may be lagging behind developments in AI technology;

- **regulatory burden** - it is important that any AI regulatory framework balances proportionate risk management with encouragement of innovation;

    - a further potential regulatory burden would be conflicting or duplicating AI regulatory requirements;

- **culture** - some firms may be less innovative than others. For example, some insurers have lost ground to Fintech firms;

    - fear of the unknown, including such as cyber threats, potential data leaks by giving own data away to an external AI-provider;

    - another fear is of ethical issues, owing to misunderstanding of the nature of such issues and how AI interacts with them;

- **market consolidation** - legacy portfolios and processes need to be integrated before improving them using AI;

    - non-cloud-based legacy systems typically take longer to integrate AI technology;

    - the insurance sector is often still heavily reliant on Excel spreadsheets;

- **explainability and interpretability concerns** - 'Black-box' AI models lack inherent transparency, making regulatory approval difficult;

- **variability of output** - GenAI models outputs are by nature variable and from time to time include hallucinations;

    - significant effort can be spent on techniques like Retrieval Augmented Generation (RAG), prompt engineering and on appropriate safeguards and screens to reduce inappropriate output, and to 'ground' the output.

Question 2: Self-reinforcing dynamics: What evidence do you see of 'winner takes most' dynamics in AI, such as data feedback loops, economies of scale, or network effects, that could entrench market positions? Conversely, could AI reduce switching costs and increase competition? Please distinguish between dynamics you observe today and those you anticipate.

38. We think the potential market transformation from AI is analogous to the disruption caused by previous step changes such as digitisation or aggregator sites. Specifically, AI has the potential to reduce friction, so permitting easier switching and increasing competition. In terms of timing, we do not see evidence of "winner takes most" from AI yet. This may take a few years to emerge.

Question 3: Control of the customer relationship: Who do you expect will control the primary customer relationship by 2030 onwards: incumbent FS firms, Big Tech, specialist AI intermediaries, or consumers' own AI agents? Do you see parallels with mobile wallets, where value is captured without becoming a traditional regulated provider? What would this shift mean for customers and for competition?

39. We are not aware of any seismic shifts in the nature of the customer relationship at present. There have been tentative moves in this direction by Big Tech firms (such as Amazon's interest in the UK insurance market around 2022) but nothing meaningful to date.

Question 4: Regulatory perimeter: 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?

40. To a limited extent this is already happening: for example, any consumer can ask a chatbot to recommend a particular insurance provider or product and will get answers. Indeed, many firms are monitoring (and trying to influence) what LLMs say about their brand, because consumers will be using them in this way. It might be difficult for the FCA to regulate such activity, beyond consumer education. However, we do not expect chatbots to supersede advised sales providers, certainly not by 2030.

Theme 3: Future Consumer Trends

Question 1: Benefits and risks: How might consumers benefit from AI-enabled retail finance from 2030 and what do you foresee as the greatest risks for consumers?

Benefits

41. AI's potential upside across retail financial services is wide-ranging. We highlight the following benefits, already being realised:

    - greater accuracy in risk assessment due to greater availability of data. For some consumers assessed as having a 'lower risk' profile, this could mean lower insurance premiums;
        - the potential here needs however to be balanced against retaining an 'appropriate' degree of pooling of risk;
        - what is appropriate pooling of risk is a key societal question;
    - insurance and wider financial service products becoming more of an holistic service, e.g. by motivating people to take more exercise or eat more healthily via fitness trackers, with the incentive of a reducing insurance premium;
    - greater data cleansing and hence data quality/ transparency. If understanding of data improves with better data quality, this could help reduce bias in data, which should ultimately benefit consumers;
    - greater affordability and accessibility of financial advice by making processes more efficient, personalised and compliance-driven;
    - AI-powered tools could provide scope for better and clearer communication with consumers, enhancing their understanding and enabling them to engage more effectively with their long-term pension savings and investments;
        - for example, personalized financial planning driven by AI insights (including account management, expense tracking and credit score management);
        - AI-powered tools could also help detect financial fraud.
    - AI tools could potentially help identify, and call out, bias in a less subjective way, and more quickly than if attempted by humans.

42. As technology advances, it is plausible that the degree/ value of beneficial impact may slow down. An AI algorithm may be akin to a 'gold mine' at outset, but successive developments may offer refinement rather than revolution.

Risks

43. AI could potentially be associated with discriminatory decisions, including in respect of individuals with vulnerabilities or with protected characteristics. This could be due to historic, inaccurate or unrepresentative datasets: flawed 'training' data could reinforce discriminatory practices. However, it is important to note that insurers have been familiar with this issue and managing it for some decades; it is an issue that pre-dates the rise of AI.

44. There is a risk of AI replicating existing biases; it may also be more biased through indirect discrimination where it is able to derive a protected characteristic from other features, for example using post code as a proxy for ethnicity, height as a proxy for gender or determining gender from driving behaviour.

45. Bias/ discrimination could lead to consumer detriment generally, and not just in relation to individuals with vulnerabilities or protected characteristics. Firms could invest in AI technology that promises to increase profits but may lose sight of the risk of greater financial exclusion. Even if such bias were inadvertent, it does not lessen any consumer harm

46. AI can increase both a firm and its customer's exposure to cyber risk. AI-driven cyberattacks can exploit vulnerabilities at scale and amplify potential impacts through the interconnectedness of systems and devices. This could be mitigated by upgrading cyber resilience standards, enhancing AI-driven cybersecurity measures and appropriate training to users of AI tools.

47. There are also cyber risks arising from using open-source AI models which can have hidden vulnerabilities or malicious code.

48. A number of current data protection concerns may become more pressing as AI becomes ever-embedded including:

    - consumers may not be fully aware of the type and extent of personal data used in product pricing processes;
    - AI may be used to launch more sophisticated cyberattacks;
    - data retention - data required to back AI-powered systems must be backed by clear data governance policies;
    - care should be taken to ensure that data is deleted securely and that it is not retained longer than regulation requires.

Question 2: Inclusion versus exclusion: Which consumer segments might 'win' or 'lose' in this new world of AI-enabled retail finance?

49. Consumers could 'win' where they enjoy the benefits under any of the scenarios set out in theme 3 question 1 above. For example, young drivers with an uncharacteristically low risk profile could benefit from more affordable motor insurance.

50. Consumers could lose out if an AI-driven risk assessment led to certain consumers being regarded as 'high risk'; this is a potential downside of greater accuracy in risk assessment. A 'high risk' assessment could lead to expensive or unaffordable financial service products, increasing financial exclusion for those affected. Other benefits need not have winners and losers however - for example, greater personalisation of products, or a more holistic service offering should largely be positive developments for most consumers.

51. There is a risk that the benefits of AI could be shared unequally between different groups within society. For example, firms seeking continued growth might prioritise younger, more technically sophisticated customers at the expense of older customers who tend to favour traditional products.

Question 3: Changes to products and services: 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.

52. As we have noted above, AI-powered tools could improve the way in which consumers engage with their long-term pension savings and investments. They could make the delivery of regulated financial advice more affordable and accessible by making processes more efficient, personalised and compliance-driven. Such tools could also help detect fraud and anomalies in financial transactions. However, greater personalisation could lead to a further shift away from the main risk pooling purpose of insurance.

53. It is important for firms to be aware that, whatever the technology used, they remain accountable for their decisions. In retail finance, consumers will expect firms to ensure transparency of decision making, so that they can explain to the customer why a particular decision has been taken.

Question 4: Agency and understanding: With the balance shifting between consumer agency and delegation to AI, how might this affect consumer understanding, financial literacy and vulnerability?

54. Vulnerable consumers need not necessarily miss out on the potential benefits listed earlier. However, AI could also offer some specific benefits to vulnerable customers. For example:

- AI could be used to help identify vulnerabilities by analysing data and identifying patterns on a real-time basis. This could help improve customer service and compliance;
    - it should however be recognised when reliance on AI is appropriate here;
        - flagging potential vulnerability is helpful;
        - subsequently relying on a chatbot to engage with the vulnerable customer may be less appropriate;
        - voice recognition software can be counter-productive in the context of vulnerable customers. Where such a customer has an 'atypical' voice, this might erroneously flag fraudulent activity;
- AI could also improve engagement with consumers, depending on their vulnerability;
    - for example, by using more accessible language, or engagement suited to specific neuro-diverse conditions;
- although a general benefit of AI is early detection of fraud and financial exploitation (as above), this may be particularly important in relation to vulnerable consumers if they are more susceptible to exploitation.
- we are aware that some are already using chatbots for insurance-related questions, with the risk that the chatbots will hallucinate answers, thus giving bad advice. (An insurer in Spain has recently been approved to sell insurance through ChatGPT.[3]) In this environment, we think it is important for the FCA to continue to prioritise initiatives that will improve consumers' financial literacy and their understanding of such risks.

Question 5: Fraud: 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?

55. We have noted potential benefits of fraud detection by AI. On the other hand, AI-driven fraud such as deepfakes are an established phenomenon, and unfortunately it seems clear that any new technology will be adopted by fraudsters very quickly to commit new types of fraud. Agentic AI in particular makes this easier for fraudsters to achieve.

[3] https://www.lifeinsuranceinternational.com/analyst-comment/openai-approves-first-insurer-app/?cf-view

Question 6: Trust: What might help make AI-driven decisions more understandable and trusted by customers, including how the use of AI may be monetised

56. Consumers could be impacted by automated decisions they do not fully understand or cannot readily challenge, which could negatively impact consumer trust. There is a range of potential solutions to improve the explainability/ 'understandability' of AI-driven decisions, including:

- the need for firms to undertake robust testing of AI models, (including comparison with alternative models);
- a requirement for transparency in AI-driven decisions, including clear explainability to consumers;
- minimum standards and disclosure requirements in communicating AI decisions to consumers;
- clear redress mechanisms for AI-driven decisions; robust compliance measures to help prevent mis-selling and scams targeting vulnerable consumers.
- transparency about ways in which involving AI could be generating revenues for others, such as:
    - using AI to profile customers to sell them products
    - cross selling or nudging based on behavioural predictions
    - monetising insights derived from consumer data
- charging fees for AI driven personalised tools

Theme 4: Future Regulatory Approach

Question 1: Outcomes-based regulation: What are the opportunities and challenges for the FCA in ensuring an outcomes-based approach to retail regulation in an AI-enabled FS industry?

57. We favour a broad principles-based, rather than rules-based approach to AI regulation. A key challenge is simply the (great) pace of evolution of AI technology. It is plausible that increasingly fewer people will be familiar with the latest technology as it emerges. The pace of evolution also means that a dynamic approach to AI regulation is necessary - for example the regulator should be able to make quick interventions where specific harms arise from use of certain technologies – but taking a principles-based approach should be better able to keep up with the pace of AI development.

58. In our view the UK's existing conduct and prudential regulatory framework is sufficiently broad that it could be adapted to encompass AI development - if a principles-based approach is adopted. Such an approach would also be key in any associated regulatory definition of AI: an overly specific definition would run the risk of becoming quickly out of date; it could also create potential loopholes.

59. One necessary task in developing the UK's current regulatory framework to be fit for purpose for AI will be to consider potential and material gaps in AI regulation. An advantage to developing the existing framework is that it should help avoid conflicting or duplicate regulatory requirements.

60. There is a balance to be struck in terms of specific AI-regulation, against adjustment of existing regulations to an AI-world. It is important to focus on how AI tools are used in different contexts, for example in relation to regulation relating to consumer fairness or financial inclusion. Principles-based and outcomes-focussed regulation may be the most appropriate way to take the opportunity of cost savings from AI automation in order to benefit consumers, increase competition and manage risks.

61. It is also important that any AI regulatory framework balances proportionate management of risk with encouraging innovation. Building and maintaining firm and consumer trust in AI is crucial.

Question 2: Regulatory levers: 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?

62. As mentioned above, we believe that the UK's existing conduct and prudential regulatory framework is sufficiently broad that it could be adapted to encompass AI development. We recognise that the regulatory levers listed have a part to play in AI regulation.

63. Any additional requirements specific to AI should complement and build on the FCA's Consumer Duty requirements - if appropriate – rather than add a separate conflicting set of requirements.

64. We note that the Review explains that the FCA will assess how relevant senior managers under SM&CR can continue to discharge their responsibilities for the deployment and maintenance of AI systems, and how these responsibilities might need to evolve. Prescribing responsibility for AI to a designated Senior Manager may not be straightforward, as it may be difficult for the relevant individual to have a span of control that would encompass all areas where AI may be in use at the relevant firm. In an insurance context, this span could be quite expansive, including inter-alia, sales, marketing, claims, underwriting and pricing functions.

65. An alternative to prescribed responsibility would be guidelines for senior managers to help them ensure that AI risks are considered; this is particularly important given the rate of evolution of AI technology.

66. Financial service firm governance frameworks are relevant to AI regulation. For many firms, we would expect their governance structures could be readily adapted to include AI. However, we would also expect firms - particularly within their second and third line of review teams - to ensure their AI governance was adequate before any material extension in their use of AI. Smaller firms may benefit from greater support in ensuring their governance arrangements were appropriate.

67. Understanding the application of AI - including ethical issues and its impact on the firm's customers - is important in avoiding conduct risk and other unintended consequences of firms' use of AI. A firm's AI governance may be augmented through the requirement for relevant members of a firm's staff to have appropriate knowledge, skills and training when embarking on AI projects.

68. Moreover, data governance is an important foundation for future modelling activity, supporting trust in the use of AI. Increased use of AI requires robust data security to protect customer information, to address data privacy concerns arising from both retail financial service firms and their customers. Firms will need a comprehensive data governance framework with corresponding controls, to protect consumer data, and also reduce the risk of bias

69. FCA review and challenge of a firm's data governance framework will be a particularly important lever in AI regulation.

70. Beyond data, a firm's governance of the AI models themselves is also important. One risk is that an AI model is so complex that it becomes opaque to model users/ owners. This then impacts consumers in relation to the explainability of decisions made by the model;

71. There is a risk that firms use AI outside their 'zone of competence' and inadvertently open themselves up to a range of unintended consequences;

    - this is plausible given the rate of technological progress in AI, or where there is a dependency on external parties in procuring AI infrastructure;

    - this risk is relevant to consumers if there is an unintended adverse impact on them.

Question 3: Supervisory and enforcement approach: 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?

72. We understand the FCA and PRA are already using ML to help them better analyse the large data sets in firms' regulatory reporting, which also increases the amount of data they request from firms. If not already doing so, we think it would be reasonable for the FCA to consider the use of GenAI to help it review and summarise unstructured data such as company report and accounts. In its use of AI, the FCA should follow the recommendations we have mentioned above, such as being transparent and aware of the limitations.

73. In terms of its wider approach to supervision and enforcement, as mentioned earlier it will be important for the FCA to find a balance between gaining benefits from cost savings and competition while maintaining sufficient controls to contain the risks of customer detriment.

Question 4: Growth: In what ways can the FCA continue to support growth and competitiveness in an AI-driven financial services industry in the future?

74. It is important that the FCA (and other regulators) and Government consider responsible use of AI, however this should be balanced with encouraging, rather than stifling innovation. Minimising consumer

detriment and ensuring financial services products are aligned to consumers' needs are important, but fostering innovation in AI is also in the wider public interest.

Question 5: Frameworks for inspiration: 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?

75. AI and hence AI regulation are global in nature. There is clearly merit in comparing/ contrasting with draft or implemented regulation of AI elsewhere, in developing a regulatory framework for AI in the UK. However, in our view it may be too early to be clear on what has worked well and where; it could be argued that globally we are on a journey in developing proportionate AI regulation.

76. Furthermore, it should be recognised that current geopolitical headwinds may have a particular bearing on global AI/ AI regulation in the short term.

77. To date however, some jurisdictions have moved faster than others. The Monetary Authority of Singapore has published principles on Fairness, Ethics, Accountability and Transparency; these principles aim to guide responsible use of AI and data analytics.

78. The EU's AI Act is the most comprehensive attempt at horizontal AI regulation to date, entering into force in 2024 with phased obligations over several years. The OECD has also done significant work on AI regulation.

79. For some firms with global operations, they may have a preference for AI regulation that is generally aligned to EU, US (or other relevant jurisdiction) equivalents.

80. It would be unfortunate if disproportionate AI regulation put the UK at a disadvantage; it is important there is adherence to global standards to ensure a level playing field.

Closing Points

81. More generally, the actuarial skillset - including critical thinking, judgement, risk management and understanding model limitations - is particularly relevant to the development of AI, not only within the financial services sector, but also right across society. We believe the IFoA, and actuaries more generally have an important role to play in the debate on the future evolution and regulation of AI.

82. Collaboration between relevant professions and disciplines will also be key, as AI has a broad reach both within individual firms and across their respective sectors. Given the UK's strength in AI research, multilateral engagement spanning academia, industry, Government and regulators has a role to play. In addition, collaboration will need to have an international dimension, both to learn from other jurisdictions and to avoid regulatory misalignment on AI.

Should you want to discuss any of the points raised please contact me, Technical Policy Manager (steven.graham@actuaries.org.uk) in the first instance.

Yours Sincerely,

Steven Graham
On behalf of Institute and Faculty of Actuaries