

Response to consultation

BANK OF ENGLAND DISCUSSION PAPER 5/22 & FCA DISCUSSION PAPER 22/4: ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

About the Investment Association

The IA champions UK investment management, supporting British savers, investors, and businesses. Our 250 members range from smaller, specialist UK firms to European and global investment managers with a UK base and manage £10.0 trillion of assets. The investment management industry supports 122,000 jobs across the UK. Our mission is to make investment better. Better for clients, so they achieve their financial goals. Better for companies, so they get the capital they need to grow. And better for the economy, so everyone prospers.

Our purpose is to ensure investment managers are in the best possible position to:

- Build people's resilience to financial adversity
- Help people achieve their financial aspirations
- Enable people to maintain a decent standard of living as they grow older
- Contribute to economic growth through the efficient allocation of capital

The money our members manage is in a wide variety of investment vehicles including authorised investment funds, pension funds and stocks and shares ISAs. The UK is the second largest investment management centre in the world, after the US and manages over a third (37%) of all assets managed in Europe.

Executive summary

For the avoidance of doubt, this paper has been prepared by a human rather than AI. We suspect that in future a clarification statement such as this may become more common.

The World Economic Forum and Cambridge University (2020)¹ found that innovative technologies such as artificial intelligence and machine learning will have either a 'very high' or 'high' long-term impact on investment returns and investment operations.

¹ ["Transforming Paradigms: A Global AI in Financial Services Survey,"](#) World Economic Forum, January 2020.

The Investment Association

Camomile Court, 23 Camomile Street, London, EC3A 7LL
www.theia.org

© The Investment Association 2023. All rights reserved.
No reproduction without permission of The Investment Association

Policy, Strategy & Innovation

 @InvAssoc  @The Investment Association

In that context, the importance of the regulatory approach is clear, both in terms of providing clarity to the industry on the framework under which AI can be deployed and also in the potential to impede innovative change that would otherwise benefit consumers.

Our response attempts to contribute to the discussion in the following areas:

- The potential benefits of introducing a sectorial regulatory definition of AI would not outweigh the potential drawbacks.
- The supervisory authorities should prioritise areas in which there are benefits for consumers and in increasing the relative competitiveness of the UK.
- Any future regulatory approach should be proportionate to the level of risk associated with the way in which these technologies are utilised and governed.
- International alignment on definitions of AI, the application of regulatory frameworks and in fostering innovation across borders will be critical.
- Focusing on the proper allocation of oversight responsibility through the Statements of Responsibility for particular SMFs is preferable to either creating a dedicated SMF role for AI or making AI a prescribed responsibility.

There are a number of questions for which we have not provided responses, either because we are not in a position to input, or because it may be too early to comment on the specifics in what is a fast moving area. In due course, we should be able to give input on areas such as desirable regulatory changes and SM&CR rules, and we look forward to contributing further as the discussion continues.

Responses to selected questions

Q1: Would a sectoral regulatory definition of AI, included in the supervisory authorities' rulebooks to underpin specific rules and regulatory requirements, help UK financial services firms adopt AI safely and responsibly? If so, what should the definition be?

We are not convinced, at this stage at least, that the potential benefits of introducing a sectorial regulatory definition of AI would outweigh the potential drawbacks. As the DP acknowledges, this is a rapidly developing field. Any sectoral definition might require periodic updating, which is itself undesirable.

Moreover, introducing sectoral definitions of AI will create unwelcome complexity. Organisations operating in multiple sectors may find themselves needing to keep track of several definitions rather than a single broad one. Similarly, as definitions proliferate across jurisdictions, international organisations will be required to navigate a crowded landscape of vying definitions.

If a sectoral definition is to be created, it is important that any definitions created for the UK market avoid contradicting definitions used for similar purposes elsewhere in other jurisdictions. This will facilitate the global regulatory coordination that is necessary to ensure the correct supervision of this emerging technology.

Q2: Are there equally effective approaches to support the safe and responsible adoption of AI that do not rely on a definition? If so, what are they and which approaches are most suitable for UK financial services?

We consider that it is appropriate to continue, as has been the case up until now, relying on a broad definition of AI. In our view the emphasis should be on the risks presented by the AI use case in question, rather than on whether it meets any number of more narrow definitions.

In this way, supervisory attention can be focused on the applications and use cases which present the greatest risks, and the definition relied upon will be flexible, agile and future proof.

Q3: Which potential benefits and risks should supervisory authorities prioritise?

Amongst other technological innovations, our member firms have attested that the use of artificial intelligence and machine learning in financial services may present profound opportunities for the investment management industry. As per the [Bank of England and FCA's survey](#) into the state of machine learning in UK financial services, the use of these technologies is already significantly widespread across critical business areas and the appetite for further expansion continues to increase.

The IA agrees that the adoption of artificial intelligence may produce significant benefits for the investment management industry and its customers, and recognises that a number of these benefits have correctly been identified in previous work undertaken by the supervisory authorities, including the ability of firms to improve operational resilience, reduce market inefficiency, and harness innovation. In this respect, it is worth emphasising that while much attention is focused on the more transformative aspects of what AI might achieve, the use cases also focus heavily on addressing everyday inefficiencies and errors that others reduce effectiveness and drive up cost.

Nonetheless, the IA believes that the supervisory authorities should prioritise the benefits for consumers that are associated with the ability of firms to provide better products and services, and in increasing the relative competitiveness of the UK.

One key area may be the ability of firms to deliver tailored services for customers. The IA has previously welcomed the review, announced by the FCA last year, of the advice-guidance boundary and the potential for firms to be able to provide guidance services to consumers via innovative channels in the future. While robo-advice has not yet realised its early potential, the ability of artificial intelligence to harness the power of large volumes of data and identify characteristics about consumers and their preferences means that interventions from firms in encouraging consumers to take limited actions where certain conditions arise may be a clear use case for these technologies. Specifically, consumers may benefit from being prompted to review the balance of their portfolio between different types of assets or putting their uninvested cash to better use. This could in turn assist the FCA's ambition to reduce the numbers of consumers with high cash balances and meet the consumer investment strategy targets of the supervisory authorities, as well as increase the capacity of investors to achieve their financial objectives.

Furthermore, the effective integration and regulation of artificial intelligence and machine learning may also increase the international competitiveness of UK-domiciled firms. As identified in this DP, when these technologies are utilised effectively in maximising their capacity to analyse large volumes of data, capture non-linear relationships, and utilise their predictive powers, this can help financial services firms create better decision-making tools, deliver optimal products and services for consumers, and provide more personalised products and services for consumers at greater scale and at a lower cost. If these technologies do indeed deliver better financial solutions and/or an improved experience for customers, then this may improve the relative attractiveness of UK-domiciled funds.

Despite these benefits, the IA also recognises that there are potential risks to the investment management industry associated with these innovative technologies. In particular, we recognise that supervisory authorities will need to be mindful of enabling an overreliance on automation, insufficient oversight or a lack of 'human in the loop' arrangements that may otherwise result in unwanted outcomes. Specifically, the IA recommends that the supervisory authorities prioritise protecting against any risks associated with these technologies interacting with market participants *in the absence of* sufficient human supervision, as well as the potential amplification of procyclical behaviour due to the increasing use of algorithmic and/or sentiment analysis across the industry.

However, and as reflected in section 3.8 of this DP, the IA believes that it is possible for supervisory authorities and FMIs alike to take the necessary steps to mitigate such risks. In relation to the risks identified above, a common driving force is the **absence of human supervision** and **adequate governance**. Permitting that FMIs are able to combine high-quality data, appropriate model choices, and good governance, this can reduce the prevalence of the risks outlined above and result in the effective *and* safe use of these technologies.

To ensure such, the IA proposes that any future regulatory approach should therefore be proportionate to the level of risk associated with the way in which these technologies are utilised and governed. More specifically, in instances where a firm has sufficient human governance over the model, then the level of regulatory scrutiny that the model is subjected to ought to be lower than would be the case with a fully automated one. Specifically, the IA recommends that in the case of AI autonomously driving trading decisions and/or providing investment guidance, it may be proper for the relevant business line SMF (e.g., SMF3 or SMF21) to have personal responsibility for ensuring the proper oversight of AI developed within the business for which they are accountable. Permitting that these discussions consider the potential risks that the use of AI applications specifically may present to the firm and its customers, then this approach may help to reduce some of the potential risks without stifling innovation and the realisation of the vast potential benefits that these technologies can deliver.

Q4: How are the benefits and risks likely to change as the technology evolves?

The IA acknowledges that some of the benefits and risks associated with artificial intelligence and machine learning will likely evolve in tandem with the evolution of the technology itself. These changes may manifest as adaptations of the benefits and risks already identified in this DP, or indeed as entirely new ones that have yet to materialise or be discussed. Nonetheless, the IA recommends that FMIs and regulators exercise caution when prophesising precisely *how* these changes will materialise and what this may mean for the financial services industry given that technological innovation is often disruptive, fast-paced, and unforeseen.

To combat the unpredictability of changing risk profiles, the IA strongly recommends that regulators position themselves so that their agility is equal to the speed at which these changes typically manifest. Without ensuring this, the extent to which the UK will be able to fulfil the objective of establishing itself as the home of a well-regulated, technologically advanced financial system, as outlined by the Economic Secretary to the Treasury (2023), will be limited.

Despite the uncertainty outlined above, our members have previously identified that one particular benefit is likely to emerge as the technology evolves. As per research previously conducted by the World Economic Forum (2020)², innovative technologies such as artificial intelligence and machine learning will have either a ‘high’ or ‘very high’ long-term impact on investment returns and investment operations. As its adoption and prevalence increases in the financial services industry, our members foresee that this will lead to greater public understanding and confidence in the use of the technology. As the familiarity and confidence of customers with this technology increases, this will increase the business case for greater adoption across firms.

All other things equal, this will likely produce two key benefits. First, it will facilitate the maximisation of the foreseen benefits of this technology beyond what has been realised to date. Second, assuming that firms are indeed successful in their attempts to utilise these technologies to deliver optimal products and investment solutions for customers, then greater trust in these outputs will also lead to greater take-up. In turn, this may reduce market inefficiencies and increase the capacity of investors to achieve their financial objectives.

Q5: Are there any novel challenges specific to the use of AI within financial services that are not covered in this DP?

The IA agrees that a number of the risks associated with artificial intelligence have correctly been identified in this DP. However, the DP focusses entirely on the risks posed by incomplete and/or inaccurate data, bias, poor model design, and insufficient governance by firms operating *within* the industry. Whilst the supervisory authorities are correct to identify these challenges, the IA would like to underline that they are rarely the product of the deliberate action(s) taken by firms. However, the same cannot be said for malevolent actors

² Ibid

operating *outside* of the financial services industry whose objective is to deliberately cause harm to market stability, firms, and/or consumers. The proliferation of powerful and widely available AI capabilities will inevitably be seized upon by criminals who can use these tools to more effectively commit fraud or cyber-attacks. The potential harm arising from such actions could arguably represent a greater threat to the regulators' supervisory objectives, particularly consumer protection, than that posed by well-intentioned firms developing and using AI in good faith.

In the case of fraud, to give one example, language models capable of generating human-like text can be used to easily produce highly convincing prose with which to manipulate victims. Many of the numerous phishing and impersonation attacks currently targeted against the industry and consumers are quite unsophisticated and therefore easy to defeat. That may not remain the case for very long.

While many AI tools include safeguards designed to prevent them being used for nefarious deeds, in many cases these can be circumvented, and cannot be relied upon completely.

Whilst the challenge of cyber-attacks has not been addressed in this DP specifically, the Belfer Center at Harvard University has published extensive research (2019)³ examining the novel specific to the use of AI by malevolent actors outside of the industry. These challenges are likely to manifest in one of two ways:

- a. **Input Attacks:** manipulating what is fed into the AI system in order to alter the output of the system to serve the attacker's goal, such as inputting new lines of code that may reduce market stability and trust in targeted FMIs. These attacks are particularly dangerous because the attack patterns do not have to be noticeable and may be completely undetectable.
- b. **Poisoning Attacks:** corrupting the process during which the AI system is created so that the resulting system malfunctions in a way desired by the attacker, such as destroying the underlying code so to prevent the operationalisation of important business services. Attacks of this nature typically target the underlying datasets; algorithms used to learn the model; and model itself.

By successfully exploiting model inputs, attackers can successfully destroy the performance of an otherwise robust model, largely irrespective of the steps taken by the firm themselves. Furthermore, due to elements of opaqueness in advanced models, such as those used in deep neural networks (DNNs), this can limit the capacity of firms to distinguish between poor model performance and being compromised by malicious actors. As a result, these potential vulnerabilities are not "bugs" that can be patched or corrected, as is done with traditional cybersecurity vulnerabilities, which may reduce the capacity of firms to mitigate such attacks.

Accordingly, the IA recommends that supervisory authorities continue to collaborate with, and leverage the expertise of, cyber resilience professionals at the National Cyber Security Centre (NCSC) to provide actionable guidance outlining the risks and mitigation strategies that firms and FMIs can implement to protect themselves against such risks.

Q8: Are there any other legal requirements or guidance that you consider to be relevant to AI?

The DP rightly refers to the EU's proposed AI regulation in the context of similar developments internationally. However, we would like to emphasise the extraterritorial reach of those proposals and the impact they will have on how global organisations approach the development of AI.

The EU's proposals will apply to "*providers placing on the market or putting into service AI systems in the Union, irrespective of whether those providers are established within the Union or a third country*", and, "*providers and users of AI systems that are located in a third country, where the output produced by the system is used in the Union*"⁴. In practice, this will likely mean that global financial services firms operating in the UK will set out to ensure that any AI applications they develop will comply with the EU rules, or any

³ ["Attacking Artificial Intelligence: AI's Security Vulnerability and What Policymakers Can Do About It" Harvard Kennedy School, August 2019](#)

⁴ European Commission, Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, Article 2 - <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52021PC0206&from=EN>

other international rules for that matter, to avoid having to develop different capabilities for use in different markets.

Q9: Are there any regulatory barriers to the safe and responsible adoption of AI in UK financial services that the supervisory authorities should be aware of, particularly in relation to rules and guidance for which the supervisory authorities have primary responsibility?

Article 22 of UK GDPR, which provides safeguards around the use of automated individual decision making and grants data subjects certain rights, has been viewed by some to be a regulatory barrier that in practice amounts to a general prohibition on automated individual decision making⁵. We note that the Data Protection and Digital Information Bill is currently making its way through Parliament. The Bill, as drafted, will amend Article 22 to clarify when it applies, and also align it with the government's AI strategy.

We welcome this development and consider it will pave the way for beneficial automated individual decision-making whilst maintaining appropriate safeguards and a route to challenge automated decisions for data subjects.

There is also an ongoing debate over the extent to which the principle of AI explainability will remain feasible in practice as models grow in sophistication. One school of thought suggests that models will become ever harder to explain, and within time firms will be forced to deploy simpler models for which they can satisfy explainability requirements, over more powerful but harder to explain ones. This might stifle innovation. On the other hand, some hold that as the sophistication of models develops, so too will the techniques for explaining them.

It is difficult to foresee how this area will ultimately develop and we are not in a position to make any bold predictions. We agree entirely with the importance of explainability, however, we do suggest that the regulatory authorities are mindful of the potential scenario whereby the desire to maintain the principle of explainability begins to compete with the ambition to drive further innovation. By implication, explainability requirements should be kept on an agile foundation, whereby the requirements can be modified in a timely manner (e.g., without having to amend legislation) as appropriate in light of future technical developments.

Q10: How could current regulation be clarified with respect to AI?

A concern expressed by some of our members is that the ambiguity surrounding how AI is currently regulated means that firms are sometimes unsure what they can and cannot pursue with regards to AI and ML.

Chapter 4 of the DP is a helpful first step in summarising existing legal requirements and the various components of the current regulatory framework. We would encourage the regulators to create new guidance in order to bring further clarity for market participants. The Kalifa Review of UK Fintech made two sub-recommendations which may be a good place to start:

"a. There should be specific guidance about the application of the PRA and FCA rules. There is a lack of clarity over how the existing rules should be applied in the context of AI, particularly in relation to the following issues, as identified by our contributors:

- *governance and accountability;*
- *"explainability" and customer understanding – i.e., how firms explain to their customers how decisions were reached and to what extent AI was involved; and*
- *the extent of human oversight or intervention required in the decision process.*

⁵ House of Commons Library, Research Briefing, Data Protection and Digital Information Bill, pp: 24-27, <https://researchbriefings.files.parliament.uk/documents/CBP-9606/CBP-9606.pdf>

b. There should be additional guidance to clarify the application of the Equality Act and the Data Protection Act, including how such legislation interrelates with relevant sector-specific regulation.”⁶

Q12: Are existing firm governance structures sufficient to encompass AI, and if not, how could they be changed or adapted?

We agree that firms will need to be mindful that an overreliance on automation, insufficient oversight and lack of ‘human in the loop’ can naturally derive unwanted outcomes and risk of harm. In addition to human oversight at the time of decision making, various aspects of the conceptual framework applicable to risk management and governance for different risk categories also apply to AI. Organisations should decide whether to treat AI as a new standalone risk category or to include it within the existing risk profile of each business process that uses AI.

With this said, assessing where firms use AI in their organisation could be the first step. Second is determining the functionality, including the source of data, transparency and explainability of the system; the limits and boundaries set for the use of AI; and the apportionment of contractual and tortious liability between programmers, suppliers, the company, and its clients in standard terms of business.

Regardless of the approach taken, the IA (2020), in collaboration with EY and Clifford Chance produced an in-depth report⁷ that outlines a number of essential mitigation measures that regulators and firms alike may wish to implement so to ensure the effective use of these technologies:

1. Monitoring exposure to AI, firms should keep track which areas of their business is exposed to AI in order to have a complete overview of where AI impacts decision making whether it be for internal purposes or client facing. This AI inventory can help ascertain actual performance against intended use of AI and act as a measure of oversight or governance of AI as a whole as well as risk assessed where needed.
2. Implementing it within Governance and Culture, further to the above, AI should be talked about throughout the organisation and treated as a risk within itself so that it receives the right level of attention and can act as an enabler for transparency throughout the organisation. For example, the risk of AI or the ethical use of AI could be integrated in the organisation’s leadership, including the company’s boards, general counsels, senior data, compliance, risk and policy teams overseeing AI risk management. They also need to determine and document management responsibility for the institution’s use of AI, with read across and a consistent approach taken across the organisation. AI should also be considered in the wider SMCR world (please see our response to the next question).
3. Ongoing testing and monitoring of AI solutions, far beyond the development stage of the solution, should involve a wide range of stakeholders and span multiple functions within FIs (e.g., HR, technology, operations and engineering). Depending on the use case and risk involved, controls may be intensive and ongoing, or less frequent (e.g., data set vetting and post-model calibration). The tests and controls used should adapt and evolve as the technology develops. Ongoing assurance work could include testing consistency between existing policies which contemplate the use of AI, e.g., GDPR compliance policies, human rights policies, competition policies, codes of conduct, and new product approval process guidelines.

Q13: Could creating a new Prescribed Responsibility for AI to be allocated to a Senior Management Function (SMF) be helpful to enhancing effective governance of AI, and why?

As noted in the DP, there is an existing certified function in the FCA’s Senior Management Arrangements, Systems and Controls Sourcebook (SYSC) covering the approval and oversight of algorithmic trading (as well

⁶ Kalifa Review of UK Fintech, page 28, <https://www.gov.uk/government/publications/the-kalifa-review-of-uk-fintech>

⁷ The Investment Association, AI and the Investment Management Industry, <https://www.theia.org/sites/default/files/2021-11/AI%20and%20the%20Investment%20Management%20Industry%20FINAL.pdf>

as deciding whether or not a trading algorithm is compliant with a firm's obligations). However, the definition of an 'algorithm' for this purpose is defined broadly as covering any "computer algorithm used in algorithmic trading", so potentially already captures algorithms incorporating AI whilst not extending to the use of AI for purposes other than trading.

With this in mind, the IA are supportive of the suggestion in the DP that the SMCR regime should be updated to reflect technological developments since the existing provisions of the Sourcebook was developed. The Certification Regime is designed around the certification of staff to perform "Significant Harm Functions", which typically include activity which involves, or might involve, a risk of significant harm to the firm or any of its customers. However, a catch-all certification function for AI may risk importing into the SMCR regime some of the drawbacks and potential uncertainties of the classification of high-risk AI under the AI Act and therefore focusing changes to the certification's regime to specific, identified applications of AI where there is an identified risk of significant harm would be preferable.

To elaborate, concentrating the oversight of AI into a single SMF function or a single individual may be challenging with reference to the varied applications of AI throughout regulated firms. It could, in fact, undermine one of the key purposes of SMCR in ensuring that business line and function-aligned SMFs have individual responsibility and accountability for activities carried out within their perimeter. For example, in the case of AI driving trading decisions, it may be proper for the relevant business line SMF (e.g., SMF3 or SMF21) responsible for overseeing that business to have personal responsibility for ensuring the proper oversight of AI developed within the business for which they are responsible, rather than their individual accountability being superseded by the concentration of AI accountability in a different, single individual. This would also reinforce the principle of technology-neutral regulation. With this in mind, focusing instead on the proper allocation of oversight responsibility for AI through the Statements of Responsibility for particular SMFs would seem to be preferable to either:

- creating a dedicated SMF role for AI; or
- making AI a prescribed responsibility (i.e., a specific responsibility that must be allocated to at least one senior manager), as some firms may not be users of AI at all, whilst others will be heavy users of AI.

Thinking further ahead to a scenario in which creative AI can be deployed, there may well come a point at which the AI programme itself may need to be required to justify its own decisions. This is a theoretical discussion at present, but one that we should keep in mind when designing an accountability regime. At what point do we accept that a human employee or senior manager no longer has effective control over the code, or is able to explain its decisions? The programme can be provided with creative licence to make business decisions or other significant actions. In that scenario the only governance lever in place for the relevant SMF to intervene by withdrawing the programme's licence to operate. At this stage we would submit that the fact that the AI is able to make complex decisions or otherwise provide unique outputs in the same way a human can should not justify onerous compliance requirements exceeding that that would apply to a human.

Q18: Are there approaches to AI regulation elsewhere or elements of approaches elsewhere that you think would be worth replicating in the UK to support the supervisory authorities' objectives?

We consider that there may be benefit in exploring a risk-based categorisation for AI applications whereby the level of regulatory requirements and supervisions is proportionate to the risks posed by the particular use case and the potential for harm. The regulatory authorities will no doubt be aware that such an approach is a feature of the EU's proposed AI framework.

Relevant factors for a risk-based approach in the context of investment management could include:

- The extent to which the AI operates autonomously
- The extent to which the AI interacts with the market
- The extent of human governance over the AI

- The potential for consumer harm resulting from the AI's use
- The potential for distorting market competition
- The potential impact on safety and soundness
- The potential impact on financial stability and market integrity.