Non-Confidential

Institute and Faculty of Actuaries

17 February 2023

Bank of England Discussion Paper DP5/22 Artificial Intelligence and Machine Learning

IFoA Response

The Institute and Faculty of Actuaries (IFoA) is a royal chartered, not-for-profit, professional body. We represent and regulate over 32,000 actuaries worldwide, and oversee their education at all stages of qualification and development throughout their careers.

Key points

The IFoA welcomes the opportunity to respond to this Discussion Paper. Although the use of AI may offer a range of societal benefits, it can also generate or exacerbate a range of challenges, which are very relevant to the wider public interest.

One potential challenge with a regulatory definition of AI is that it may become invalid, irrelevant or outdated quite quickly, given the ongoing rate of technological progress. A further challenge is if an AI definition adopted were too complex, specific or constraining, it could restrict the scope of corresponding regulation and may also create potential loopholes. As an alternative, we suggest devising a simple, broad principles-based description of what is meant by AI.

Although AI's potential upside is wide-ranging, we highlight the following benefits: greater accuracy in risk assessment; insurance becoming more of an holistic service; greater data cleansing and hence quality/ transparency; automation of routine tasks; AI could also help provide wide-ranging societal benefits, such as climate change risk modelling and analysis.

One key general risk is 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.

Specific key risks include: an increase in financial/ insurance exclusion - this is a potential downside of the greater accuracy in risk assessment; conduct risk – i.e. unfair treatment of customers by an algorithm because there is a bias within it; modelling risk - i.e. a model is so complex it becomes opaque to model users/ owners.

We agree that AI could be associated with discriminatory decisions, including in respect of individuals with protected characteristics. Although such discrimination could be inadvertent, this does not lessen any consumer harm. One potential mitigant of bias in models/ data would be for firms to undertake robust testing, although there may be a limit on how much testing can be accomplished every time a model is updated.

Although we have not identified any current regulatory barriers, it is important that any regulatory framework for AI balances proportionate management of risk with encouragement of innovation. One potential form of regulatory barrier would be conflicting or duplicate regulatory requirements.

We believe that the IFoA has an important role to play in the debate on the future evolution of regulation of AI and Machine Learning. We hope the Bank of England finds our response helpful and constructive, and we would be delighted to discuss it with the Bank.

Beijing	14F China World Office 1 · 1 Jianwai Avenue · Beijing · China 100004 · **Tel**: +86 (10) 6535 0248
Edinburgh	Level 2 · Exchange Crescent · 7 Conference Square · Edinburgh · EH3 8RA · **Tel**: +44 (0) 131 240 1300
Hong Kong	1803 Tower One · Lippo Centre · 89 Queensway · Hong Kong · **Tel:** +852 2147 9418
London (registered office)	7th Floor · Holborn Gate · 326-330 High Holborn · London · WC1V 7PP · **Tel**: +44 (0) 20 7632 2100
Oxford	1st Floor · Park Central · 40/41 Park End Street · Oxford · OX1 1JD · **Tel**: +44 (0) 1865 268 200
Singapore	163 Tras Street · #07-05 Lian Huat Building · Singapore 079024 · **Tel**: +65 6906 0889

www.actuaries.org.uk

1. The Institute and Faculty of Actuaries (IFoA) welcomes the opportunity to respond to the Bank of England's Discussion Paper DP5/22 on Artificial Intelligence and Machine Learning. As is recognised in the DP, although the use of AI may offer a range of societal benefits, it can also generate or exacerbate a range of challenges. The DP is helpful in setting out these benefits and challenges from a regulatory perspective, and the issues considered within are very relevant to the wider public interest.

2. In developing our response to the DP, we have drawn upon input from members working in general insurance and who are members of the IFoA's Data Science community. This IFoA activity includes leading on regulation, professionalism and ethics of data science.

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

4. Artificial Intelligence and Machine Learning have significant potential to promote innovation across a range of areas in which our members practise, including general, life and health and care insurance. However, as they delve ever deeper in our lives, Artificial Intelligence/ Machine Learning raise significant questions over ethics and the public interest. These are challenges the IFoA has been considering in an insurance context for over 150 years.

5. These are challenges the IFoA has been considering in an insurance context for over 150 years. Anticipating the importance of Artificial Intelligence/ Machine Learning for financial services, in 2016 the IFoA established a separate working group focusing on data and modelling problems. As of now this community has developed into a board consisting of four pillars (research, professionalism and ethics, engagement, lifelong learning) and accommodating more than 1,000 members. In 2019 the IFoA launched an ethical guide to data science in conjunction with the Royal Statistical Society. More recently the IFoA published non-mandatory ethical and professional guidance on data science for our members.

6. Given the above, we believe that the IFoA has an important role to play in the debate on the future evolution of regulation of Artificial Intelligence and Machine Learning. We hope the Bank of England finds our response helpful and constructive, and we would be delighted to discuss it with the Bank.

7. The DP considers both Artificial Intelligence (AI) and Machine Learning. In the paragraphs that follow, any comments we make in relation to AI would also apply in relation to Machine Learning as appropriate, unless otherwise stated.

8. We have responded to those set questions within the DP where we have specific points to raise.

Question 1: 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?

9. One potential challenge with a regulatory definition of AI is that it may become invalid, irrelevant or outdated quite quickly, given the ongoing rate of technological progress. For example, if AI were defined using a specific list of current algorithms, that could become out of date within a couple of years as new algorithms emerged. A further challenge is if an AI definition adopted were too complex, specific or constraining, it could restrict the scope of corresponding regulation. An overly-specific or constraining definition could also create potential loopholes.

10. However, clearly the regulator needs to describe what is within the scope of AI regulation, and hence we suggest devising a simple, broad principles-based description of what is meant by AI. We note that the DP refers to AI and Machine Learning and then draws a distinction to note that AI can pose novel challenges. If this distinction is to be reflected in regulation, it is important to clarify where the boundaries between any AI and Machine Learning lie in any such principles-based description.

Question 2: 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?

11. As in our response to question 1 above, a broad principles-based description of what is meant by AI may be effective, and could avoid obsolescence or a restrictive definition which could lead inadvertently to regulatory loopholes. To avoid the risk of an overly-broad description, one option would be to focus on regulatory outcomes intended, and regulatory challenges to be avoided in the use of AI. Emphasising the importance of ethical use of AI within any such description could also be useful. We note that the cross-sectoral principles tailored to AI characteristics set out in the DP may be a useful reference point for a broad definition.

Question 3: Which potential benefits and risks should supervisory authorities prioritise?

12. Although AI's potential upside is wide-ranging, we highlight the following benefits:

- greater accuracy in risk assessment (greater portfolio segmentation) due to greater availability of data; 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 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 quality and transparency. If understanding of data improves with better data quality, this could help reduce bias in data;
- automation of routine tasks, freeing data users to focus on more interesting/ value-adding problems;
- from a 'bigger picture' perspective, AI could also help provide wide-ranging societal benefits, such as climate change risk modelling and analysis for example.

13. One key general risk is 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.

14. Specific key risks include:

- an increase in financial/ insurance exclusion. This is a potential downside of the greater accuracy in risk assessment listed in the benefits above. If AI risk assessment led to certain consumers being regarded as high risk, it could lead to expensive or unaffordable financial service/ insurance products, increasing financial exclusion;
- conduct risk – i.e. unfair treatment of customers by an algorithm because there is a bias within it;
    - one notable risk here is rating by gender by proxy, such as where an algorithm does this automatically;
- modelling risk – i.e. a model is so complex it becomes opaque to model users/ owners, or significant key person dependency because only the person who built the model really knows how it works;

- over-complexity in modelling – where a model is implemented using sophisticated methods, when a simpler, more transparent approach could give similar results for lower investment.

Question 4: How are the benefits and risks likely to change as the technology evolves?

15. A key challenge is simply the (great) pace of evolution in AI/ Machine Learning technology, as mentioned in our points above on an AI definition. The pace of development could also increase the risks discussed in our response to Question 3 above. For example, it is plausible that increasingly few people will be familiar with the latest technology as it arrives. A further issue is that firms may be left with 'legacy' technical debt from models that are by no means old, but have been superseded given the great rate of technological progress.

16. As technology advances, it is plausible that the degree/ value of beneficial impact may slow down. An algorithm may be akin to a 'gold mine' at outset, but successive developments may offer refinement rather than revolution. Conversely, it could be argued the level of risk may increase as technology develops, to the extent that an increasing degree of decision making will be automated.

17. On a more specific point considering one application of AI and insurance - driverless cars - the nature of insurable risk may evolve from personal liability to covering the autonomous vehicle producer's liability.

Question 6: How could the use of AI impact groups sharing protected characteristics? Also, how can any such impacts be mitigated by either firms and/or the supervisory authorities?

18. We agree with the risk highlighted in the DP that AI could be associated with discriminatory decisions, including in respect of individuals with protected characteristics. We recognise this could be due to historical or unrepresentative datasets, although this could lead to consumer detriment generally, and not just in relation to individuals with protected characteristics. The DP notes that such discrimination could be inadvertent, however this does not lessen any consumer harm.

19. One potential mitigant of bias in models/ data would be for firms to undertake robust testing (including comparison with alternative models), although there may be a limit on how much testing can be accomplished every time a model is updated. Hence the development of proportionate industry best practice could be useful here.

Question 7: What metrics are most relevant when assessing the benefits and risks of AI in financial services, including as part of an approach that focuses on outcomes?

20. In relation to the benefits of AI in financial services, firms' metrics on expected costs savings or improved profits may be relevant. Future benefits due to time optimisation, and if measurable, less stress due to shorter employee working hours may also be worthwhile recording.

21. With respect to risks, one important area to measure would be conduct outcomes: such as complaints, treatment of customers with additional needs, and the new (FCA) Consumer Duty outcomes.

Question 9: 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?

22. Although we have not identified any current regulatory barriers, it is important that any regulatory framework for AI balances proportionate management of risk with encouragement of innovation. It

would be suboptimal if AI regulation had the effect of placing some practitioners at an advantage over others.

23. One potential form of regulatory barrier would be conflicting or duplicate regulatory requirements. We welcome the acknowledgement within the DP of overlapping rules and principles in relation to AI, and it is helpful that the DP recognises that regulation should be developed to address gaps in the existing framework. Regulatory development on this basis would hopefully minimise duplicate or conflicting requirements.

24. We also note reference to the FCA's impending new Consumer Duty within the DP. We agree this is relevant (and should make an important addition) to the wider AI regulatory framework. Given that the new Consumer Duty will start to take effect from July 2023, we suggest it would be useful to take stock of its success in identifying residual regulatory gaps. Any additional regulation specific to AI should complement and build on the new Consumer Duty requirements - if appropriate – rather than add a separate conflicting set of requirements.

Question 12: Are existing firm governance structures sufficient to encompass AI, and if not, how could they be changed or adapted?

25. 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 were adequate before any material extension in their use of AI.

26. Smaller firms may benefit from greater support in ensuring their governance arrangements were appropriate. Bank support through the use of open data - such as data.gov.uk - may be beneficial to firms in their data use/ model training. This could potentially extend to creation and release of open data.

27. More generally, 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. Such training should include an ethical dimension, and not just technical Data Science. As noted in our general comments within this response, the IFoA launched a Data Science certificate and we are currently working on the addition of more granular and technical data science content more widely within our fellowship syllabus.

Question 13: 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?

28. Prescribing responsibility for AI to a designated Senior Management Function (SMF) 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. An alternative to prescribed responsibility for SMFs 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.

29. The role of Chief Data Officer where it exists is also relevant to AI governance, as is, in the context of GDPR, Data Privacy Officers. Similar comments apply to these roles as for SMFs, in relation to considering AI governance.

Question 14: Would further guidance on how to interpret the 'reasonable steps' element of the SM&CR in an AI context be helpful?

30. Given that the concept of 'reasonable steps' is a core element of the SM&CR, guidance on interpretation in an AI context would be useful, if this could be provided.

Question 15: Are there any components of data regulation that are not sufficient to identify, manage, monitor and control the risks associated with AI models? Would there be value in a unified approach to data governance and/or risk management or improvements to the supervisory authorities' data definitions or taxonomies?

31. As noted in the DP, and in our response above, a firm's AI governance is key. 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. Moreover, ***data governance*** is an important foundation for future modelling activity, supporting trust in the use of AI.

Question 16: In relation to the risks identified in Chapter 3, is there more that the supervisory authorities can do to promote safe and beneficial innovation in AI?

32. We would reiterate a point made earlier. Whilst it is important that any regulatory framework for AI focusses on risk management (including potential consumer detriment), this should be balanced with encouraging, rather than stifling innovation. As with minimising consumer detriment and ensuring financial services products are aligned to consumers' needs, fostering innovation is also in the wider public interest.

Question 18: 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?

33. As the DP explains, there has been ongoing global debate on the regulation of AI. There is clearly merit in comparing/ contrasting with draft or implemented regulation of AI elsewhere, in developing an AI framework for 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.

34. 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.

35. In the specific context of any review of GDPR requirements, it may be worthwhile comparing experience of the GDPR with the Swiss Data Protection Act 2020.

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

Yours Sincerely,

Steven Graham
On behalf of Institute and Faculty of Actuaries