Bank of England PRA

The Rt Hon the Baroness Stowell of Beeston MBE
Chair of the Select Committee on Communications and Digital
House of Lords
London
SW1A 0PW

Sir Jon Cunliffe
Deputy Governor, Financial Stability
Bank of England

Sam Woods
Deputy Governor, Prudential Regulation and CEO of the Prudential Regulation Authority
Bank of England

5 October 2023

Dear Baroness Stowell,

We write in response to your letter, dated 14 September 2023, regarding the work and capacity of the Bank of England and Prudential Regulation Authority (PRA) relating to Large Language Models (LLMs). Please see the Annex to this letter for responses to your questions.

The Bank and PRA's objectives and technology-agnostic approach to regulation

1. The Bank's statutory objectives are to maintain monetary and financial stability. The PRA – as part of the Bank – is responsible for the prudential regulation and supervision of around 1,400 banks, building societies, credit unions, insurers and major investment firms. The Bank also regulates certain financial market infrastructures (FMIs). Given those responsibilities, the Bank and PRA's focus is on the macro-financial and prudential implications of the use of AI in financial services.

2. The Bank and PRA adopt a technology-agnostic approach to supervision and regulation. We do not regulate technologies, rather we regulate the use of technology by our regulated firms, where such use gives rise to risks to our objectives. As such, our core principles, rules and regulations do not usually mandate or prohibit specific technologies. However, certain technologies may raise novel challenges, which may mean it is difficult for firms to understand how existing

rules apply to that technology. In those cases, we may issue guidance or use other policy tools to clarify the application of our regulatory framework.[1]

Work to date on the use of AI/ML in financial services

3. The Bank and PRA – in collaboration with the Financial Conduct Authority (FCA) – have been exploring the implications of the use of AI and machine learning (ML) in financial services for a number of years. In 2019, we published the results of a joint survey to better understand the **use of ML in UK financial services**. From 2020 to 2022, the Bank and FCA ran an AI Public-Private Forum (AIPPF). The AIPPF's **final report**, published in February 2022, explored the barriers, challenges, and risks related to the use of AI in financial services and potential ways to address them. The AIPPF made clear that the private sector wanted regulators to have a role in supporting the safe adoption of AI in UK financial services, and a **further Bank-FCA survey in 2022** demonstrated that financial services firms were continuing to adopt machine learning at pace.

4. In October 2022, the Bank, PRA and FCA published a **Discussion Paper** on the use of AI and ML in financial services. It considered how key existing sectoral legal requirements and guidance in UK financial services apply to AI, and how they support consumer protection, competition, the safety and soundness of individual firms, market integrity, and financial stability and/or how any additional intervention may support the safe and responsible adoption of AI in UK financial markets. With regard to the PRA's secondary objectives, the paper also sought feedback on potential regulatory barriers and set out an overview of international regulatory developments in this area.[2]

Large Language Models and our regulatory approach

5. Since the publication of the Discussion Paper last year, advances in large language models have led to a number of concerns in the public discourse, including: the potential for them to amplify bias in the training data, their lack of interpretability[3], their capability to generate fluent but factually inaccurate outputs (so-called 'hallucinations'); the instability of their outputs to subtly different prompts; and the potential to use them for malicious purposes. A number of financial firms, as well as providers of services to the financial sector, have stated publicly that they are experimenting with the use of Large Language Models e.g. **Morgan Stanley**,

[1] The PRA Business Plan 2023-24 notes that the "PRA and other UK supervisory authorities are interested in the safe and responsible adoption of these [AI] technologies in UK financial services, including considering the appropriate role of policy and regulation".
[2] The Discussion Paper also set out how financial services regulation sits alongside a body of cross-sectoral legislation and regulation, including data protection laws, security and intellectual property.
[3] The internal structure of a Large Language Model is observable (e.g. the number of layers in the deep neural net and the value of a given parameter at a point in time) but not interpretable (i.e. it is not clear how the optimised parameters yield a sophisticated natural language output).

Bloomberg, though these public announcements – along with supervisory intelligence – indicate that regulated financial firms are exploring relatively low-risk use cases at this time. Many of the risks arising from LLMs are relevant to the use of Machine Learning in general (though they may be exacerbated in the case of a Large Language Model); some are less relevant to the Bank and PRA's objectives.

6. Given our technology-agnostic approach, we do not expect to develop a specific regulatory regime for the use of LLMs – or AI in general. Instead we will ensure that existing areas of prudential policy are sufficient to encompass firms' use of AI/ML, including LLMs. We believe that AI models are better seen on a spectrum of increasingly complex quantitative models, with requirements determined by characteristics (eg autonomy in decision-making, dynamic learning and calibration, transparency). In May this year, the PRA published **Supervisory Statement 1/23** 'Model risk management principles for banks'. The principles are designed to be sufficient to identify, manage, monitor, and control the risks associated with complex quantitative models, including AI/ML models.

7. In addition, the PRA also has clear expectations regarding Board and Senior Manager ownership of operational risks, including as associated with use of models, outsourcing and third-party risk management (e.g. **Supervisory Statements 1/21** and **2/21**). The Bank has equivalent policies for the FMIs it regulates (e.g. **Operational Resilience**, **Outsourcing and Third Party Risk Management**). There are other regulatory measures under review by the Bank and PRA, for example on **critical third parties to the UK financial sector** and the Senior Managers and Certification Regime[4], which are also likely to be relevant to firms' use of AI/ML.

8. We will continue to supervise the risks that arise from firms' using models in line with existing supervisory guidance, and as AI/ML models (including LLMs) become more prevalent and used for activities that matter for our objectives, then they will naturally begin to fall within the scope of our supervisory reviews.

9. The Bank and PRA also recognise the importance of collaborating with other UK regulators. We will continue to work closely with the FCA, other UK sectoral regulators and international counterparts on how they are approaching the regulatory issues posed by the use of AI – including LLMs – within their remits.

10. We are confident our regulatory and supervisory approach will deliver against the expectations set out in the government's March 2023 white paper on AI regulation.

We hope that this is letter is helpful to you.

---

[4] In March 2023, HM Treasury launched a Call for Evidence to review operational aspects of the SM&CR, alongside a FCA and PRA joint Discussion Paper.

Yours sincerely,

Sir Jon Cunliffe
Deputy Governor, Financial Stability

Sam Woods
Deputy Governor, Prudential Regulation, and CEO of the Prudential Regulation Authority

Annex: Responses to questions

1. How many staff do you employ specialising in AI-related governance (if any)?

11. We have interpreted this question to be about staff involved with the regulation of the use of AI by Bank and PRA-regulated firms. Our current prudential policy development (see cover letter) is led by the Bank's Fintech Hub and Prudential Policy Directorate. The purpose of the Fintech Hub is to identify fintech topics with potential implications for the Bank at an early stage, articulate what they might mean for the Bank's mission, and work with others across the Bank to shape a response to them. The use of AI/ML in financial services has been one of two priorities for the Fintech Hub for the past several years, alongside developments in cryptoasset markets and distributed ledger technology more broadly. The Fintech Hub comprises around 8FTE, of which around 1.5FTE is allocated to AI regulation.

12. In delivering this work programme, contributions from a significant number of other colleagues across the Bank and PRA have been required (as well as engagement with our international counterparts). The 2022 joint Bank-PRA-FCA Discussion Paper (and related appendices) was a particularly resource-intensive project. The publication was supported for around six months by a cross-Bank and PRA working group of around 40 colleagues, who contributed to the programme alongside their other responsibilities. This group included staff with expertise in prudential policy, data, banking supervision, insurance supervision, FMI supervision, supervisory risk specialists, advanced analytics, and legal expertise (that is in addition to the resource contributed by the FCA).

2. How many data scientists (or equivalent) do you employ (if any)?

13. At present, the Bank explicitly identifies data scientist roles for members of staff up to Heads of Division. As of 25 September there were 82 such roles, most of which were in the Monetary Policy area, which includes the Data and Analytical Transformation Directorate, and in the PRA: 37 and 30 roles respectively. Heads of Division and more senior staff may have data science expertise but their roles include a significant amount of management and strategic direction so are not included in these numbers.

14. FMI Directorate (FMID), alongside supervising FMIs in the UK, owns and analyses some of the biggest datasets in the Bank[5]. These datasets provide important insights to the Bank's policy committees: the Financial Policy Committee (FPC) and the Prudential Regulation Committee (PRC) and the Monetary Policy Committee (MPC). The datasets are also regularly used in the Financial Stability Report, research papers and speeches. To manage these datasets, FMID has around 10 dedicated data scientists. The FMID team also works and collaborates with other

[5] www.bankofengland.co.uk/financial-stability/trade-repository-data

data scientists embedded in different business areas of the Bank to facilitate a wider use. These collaborations make it possible to apply basic machine learning tools on FMID data. Correspondingly, FMID has been recently giving an increasing level of interest to upskilling data analysts and scientists in new technologies and ML applications.

15. The PRA has an ambitious programme of work to strengthen and transform its data-related capabilities. It established in 2021 a dedicated RegTech, Data & Innovation Division (RDID) to deliver a wide-ranging programme of work to bolster the PRA's efficiency, effectiveness, and data culture, through phased investment in tools, technology, processes, and skills.[6] This includes building targeted machine learning tools to support front line supervision, including for processing natural language, alongside other work. Resources in this area have ramped up over the past few years, and are currently around 35 FTE overall, of which roughly one quarter are involved in developing specialist machine learning-related tools and the associated digital skills work for all PRA staff. In addition, a number of staff within the PRA outside RDID are involved in building these capabilities, supported by enablement technologies being delivered as part of the Bank's wider Data & Analytics agenda and expertise in the use of machine learning and natural language processing in other areas of the Bank. As well as being valuable for performing supervision, this work helps PRA staff develop and maintain skills needed to understand AI tools used by firms.

3. Do you anticipate these figures changing significantly in the next three years?

16. To meet its mission in a changing world, the Bank also needs to make the best use of data, analytics and technology. Doing so means our decisions are informed with the best possible information and analysis; our operations are even more efficient and risks better managed than today; and we publish timely, reliable, and high quality data and statistics for the public good. In recognition of the importance of data and how it is used, the Bank's Independent Evaluation Office has been conducting an evaluation of how the Bank uses data to support its policy objectives. In turn, the Bank is refreshing its Data and Analytics strategy. Both the IEO evaluation and the Bank's response will be published later in 2023.

17. A key element of the refreshed strategy is to enable the Bank to take full advantage of emerging innovative technologies such as Artificial Intelligence and Machine Learning. This will include identifying the most promising use cases around the Bank, exploring technical requirements, and establishing and socialising necessary safeguards, including an ethical framework, that help the Bank use AI safely.

18. The Bank recognises that it has much to learn from others about the effective use of AI. It will therefore actively engage with external organisations to better understand data and analytics best practice, share experience, and work on joint

[6] Further details can be found in the PRA Business Plan 2022-23.

projects where appropriate. The Bank will also consider how to establish a technology environment to enable experimentation in innovative technologies and ensure that the resulting tools can be put into production.

19. We have also announced that we will be undertaking a Banking Data Review – a unique opportunity to modernise regulatory returns in support of our broader D&A agenda. We are developing the full work plan for that and expect resourcing to increase over the coming years.

20. At this stage we do not anticipate a change in resourcing related to the regulation of the use of AI by our regulated firms. As always, the Bank and PRA will reprioritise analytical and supervisory resources as warranted by developments.

4. Do you expect any new regulatory issues of concern to arise within your remit in the next three years in relation to advances in large language models?

21. We have not, at this stage, identified any new regulatory issues of concern (see paragraphs 6-9 of cover letter).

5. Do you anticipate a requirement to conduct an audit of algorithmic processes used in large language models in the next three years? (For reference we refer to the criteria set out in a paper by the Digital Regulation Cooperation Forum on algorithmic audits: governance, empirical, and technical. Please provide an indication for each of these three types).

22. We do not anticipate a introducing a specific requirement to conduct an audit of algorithmic processes used in large language models.

23. As noted in the cover letter, we believe that AI models are better seen on a spectrum of increasingly complex quantitative models, with requirements determined by characteristics. In May this year, the PRA published **Supervisory Statement 1/23** ‘Model risk management principles for banks’.[7] The principles are designed to be sufficient to identify, manage, monitor, and control the risks associated with complex quantitative models, including AI/ML models.

24. The purpose of SS1/23 is to support firms to strengthen their policies, procedures, and practices to identify, manage, and control the risks associated with the use of

[7] SS1/23 ‘Model risk management principles for banks’ is relevant to all regulated United Kingdom (UK)-incorporated banks, building societies and PRA-designated investment firms with internal model approval to calculate regulatory capital requirements. Credit unions, insurers, and reinsurers are not in scope of the MRM expectations in SS1/23.

all models, developed in-house or externally, including vendor models, and models used for financial reporting purposes.[8]

25. SS1/23 is structured around five high-level principles designed to cover all elements of the model lifecycle: Model identification and model risk classification (Principle 1), Governance (Principle 2), Model development, implementation and use (Principle 3), Independent model validation (Principle 4), and Model risk mitigants (Principle 5).

26. The model risk management principles are supported by a number of sub-principles and encompass all elements of the model lifecycle. For example, Principle 4 on independent model validation has sub-principles relating to (i) the independent validation function; (ii) the independent review of a model; (iii) process verification; (iv) model performance monitoring; and (v) periodic revalidation.

27. We will continue to supervise the risks that arise from firms' using models in line with existing supervisory guidance, and as AI/ML models (including LLMs) become more prevalent and used for activities that matter for our objectives, then they will naturally begin to fall within the scope of our supervisory reviews.

6. Notwithstanding your answer to question 5, do you believe your organisation is well positioned to conduct audits of large language models at the (a) development and (b) deployment stage, to ensure compliance with your existing duties and the expectations in the Government's white paper?

28. See answer to question 5.

7. (If applicable): Please could you provide an estimate of how many comparable audits you have undertaken in the last five years?

29. See answer to question 5.

8. In your view, what are your main powers and duties most applicable to delivering on the five principles set out in the white paper?

30. The PRA has two primary objectives: a general objective to promote the safety and soundness of the firms it regulates, focusing on the adverse effects that they can have on the stability of the UK financial system; and an objective specific to insurance firms, to contribute to securing an appropriate degree of protection for policyholders. The PRA also has two secondary objectives – to act, so far as

---

[8] The principles are designed to complement, not supersede, existing supervisory expectations that have been published for selected model types, including internal capital models (Pillar 1) for credit risk (**SS11/13 – Internal Ratings Based (IRB) approaches**) and stress testing models (**SS3/18 Model risk management principles for stress testing**).

reasonably possible, in a way which (i) facilitates effective competition in the markets for services provided by PRA-authorised firms, and, (ii) under the Financial Services and Markets Act 2023, facilitates, subject to alignment with relevant international standards, the international competitiveness of the UK economy, and its growth in the medium to long-term.

31. As noted, the Bank's statutory objectives are to maintain monetary and financial stability. The Bank is also a regulator of financial market infrastructures (FMIs). Given those responsibilities, the Bank and PRA's focus is on the broad macro-financial and prudential implications of the use of AI in financial services.

32. The Bank and FCA's joint Discussion Paper on firms' use of AI sets out in detail how our current regulatory framework relates back to these objectives in the context of firms' use of AI. As per question 4, the Bank and PRA adopt a technology-agnostic approach to supervision and regulation, focussed on prudential outcomes rather than mandating or prohibiting specific technologies.

33. The principles in the white paper most relevant to the Bank and PRA's objectives are those on 'safety, security, robustness', 'appropriate transparency and explainability' and 'accountability and governance'. On safety, security, and robustness, the Discussion Paper for instance outlines our policies for the supervision of data and model risk management, which includes appropriate consideration of transparency and explainability. Regarding accountability and governance, the Discussion Paper covers questions such as accountability for AI applications under the 'Senior Managers & Certification Regime', as well as responsibilities of the Board under PRA SS5/15 'Corporate Governance: Board Responsibilities'.

34. With regard to the PRA's secondary objectives, the Discussion Paper also sought feedback on potential regulatory barriers and set out an overview of international regulatory developments in this area, for the purposes of comparison (Appendix 2 of the DP). We are currently analysing the feedback to the Discussion Paper and considering whether further prudential policy development is needed to ensure our rules are suitable to all approaches and technologies firms may apply.

9. Are there any areas where you anticipate your existing powers and duties will fall short of what is needed to deliver against the expectations in the white paper?

35. We have not identified any particular gaps to date across the AIPPF, Discussion Paper, firm surveys, or supervisory work. The Bank and PRA supervisors – and supervisory risk specialists – will continue to monitor how regulated firms are adopting AI/ML, including Large Language Models and other forms of generative AI. As noted in the white paper, regulatory coordination will be crucial, and the Bank and PRA will continue to work with other UK regulators, as well as

international counterparts, on how they are approaching the regulatory issues posed by AI within their remits and any potential gaps.

10. How confident are you that there are sufficient co-ordination structures, independent of Government, between regulators outside the Digital Regulation Cooperation Forum in relation to regulating large language models? a. (If applicable): please provide examples of established structures that are working well.

36. The Bank and PRA are not currently involved with any regulatory coordination structures specific to the use of LLMs.

37. However, the Bank and PRA agree that regulatory collaboration is crucial to ensure the safe use of AI/ML by the financial sector, as evidenced by our extensive joint work with the FCA, including the aforementioned Machine Learning surveys in 2019 and 2022, the AIPPF and the 2022 Discussion Paper. We will continue to collaborate closely with the FCA as we consider our respective next steps, including the scope to issue joint guidance, where appropriate.

38. There are also a number of regulatory coordination structures relating to the use of AI generally. The Bank is represented on the Digital Regulation Cooperation Forum's (DRCF's) quarterly regulators' roundtable, and we are exploring further avenues for collaborating with the DRCF (e.g. the Bank recently participated in a DRCF workshop on generative AI). The Bank is also an active member of the Information Commissioner's Office's AI and Regulators Working Group, and is represented on the Alan Turing Institute's Standards Hub Regulators' Forum.