

UK financial regulators publish industry feedback on AI and machine learning

Source copy kept by wheresthe.ai for the story "UK financial regulators publish industry feedback on AI and machine learning" (<https://wheresthe.ai/d/gb-boe-pra-fca-fs2-23-ai-feedback-2023-10/>).

Source: <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/october/artificial-intelligence-and-machine-learning>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 2a365e8aa46863fa22ec3f47e1885ee0d1e34a9c8ed6b7a9bd870e74e88efb04

Documents the source offers: <https://www.mas.gov.sg/~media/MAS/News%20and%20Publications/Monographs%20and%20Information%20Papers/FEAT%20Principles%20Final.pdf>

FS2/23 – Artificial Intelligence and Machine Learning

Feedback statement 2/23

Related links

Related links

DP5/22 – Artificial Intelligence and Machine Learning

Published on 26 October 2023

Executive summary The Bank of England (the Bank) (including the Prudential Regulation Authority (PRA)), and the Financial Conduct Authority (FCA) (collectively 'the supervisory authorities') published a discussion paper (DP) 5/22 – Artificial Intelligence and Machine Learning – in October 2022 to further their understanding and to deepen dialogue on how Artificial Intelligence (AI) may affect their respective objectives for the prudential and conduct supervision of financial firms. The DP is part of the supervisory authorities' wider programme of work related to AI, including the AI Public Private Forum (AIPPF), and its final report published in February 2022. This feedback statement (FS) provides a summary of the responses to DP5/22. Its aim is to acknowledge the responses to the DP, identify themes, and provide an overall summary in an anonymised way. The FS summarises the responses to DP5/22 as far as they concern the matters raised in that DP. It does not include policy proposals, nor does it signal how the supervisory authorities are considering clarifying, designing, and/or implementing current or future regulatory proposals on this topic. DP5/22 received 54 responses from a wide range of stakeholders. The chart below shows the number of respondents by type of institution. Industry bodies accounted for almost a quarter of respondents with banks accounting for a further fifth. There was no significant divergence of opinion between sectors.

The key points made by respondents were: A regulatory definition of AI would not be useful. Many respondents pointed to the use of alternative, principles-based or risk-based approaches to the definition of AI with a focus on specific characteristics of AI or risks posed or amplified by AI. As with other evolving technologies, AI capabilities change rapidly. Regulators could respond by designing and maintaining 'live' regulatory guidance ie periodically updated guidance and examples of best practice. Ongoing industry engagement is important. Initiatives such as the AIPPF have been useful and could serve as templates for ongoing public-private engagement. Respondents considered that the regulatory landscape is complex and fragmented with respect to AI. More coordination and alignment between regulators, domestic and international, would therefore be helpful. Most respondents said that data regulation, in particular, is fragmented, and that more regulatory alignment would be useful in addressing data risks, especially those related to fairness, bias, and management of protected characteristics. A key focus of regulation and supervision should be on consumer outcomes, especially with respect to ensuring fairness and other ethical dimensions. Increasing use of third-party models and data is a concern and an area where more regulatory guidance would be helpful. Respondents also noted the relevance of DP3/22 – Operational resilience: Critical third parties to the UK financial sector . AI systems can be complex and involve many areas across the firm. Therefore, a joined-up approach across business units and functions could be helpful to mitigate AI risks. In particular, closer collaboration between data management and model risk management teams would be beneficial. While the principles proposed in CP6/22 – Model risk management principles for banks (which have since been published by the PRA as SS1/23 – Model risk management principles for banks) were considered by respondents to be sufficient to cover AI model risk, there are areas which could be strengthened or clarified to

address issues particularly relevant to models with AI characteristics. Respondents said that existing firm governance structures (and regulatory frameworks such as the Senior Managers and Certification Regime (SM&CR)) are sufficient to address AI risks.

1: Introduction 1.1 This feedback statement (FS) provides a summary of the responses to the Bank and the FCA's discussion paper DP5/22 – Artificial Intelligence and Machine Learning . The DP, published in October 2022, set out the Bank and the FCA's views on the issues surrounding use of artificial intelligence (AI) and machine learning (ML) in UK financial services. 1.2 The aim of this FS is to provide a summary of responses to the DP. This FS does not include policy proposals nor does it signal how the supervisory authorities are considering clarifying, designing, and/or implementing current or future regulatory proposals on this topic. 1.3 The responses described in this FS are presented in an anonymised way. This FS sets out broad themes from the responses to DP5/22 with the intention of providing an overall summary of the responses. Not all respondents answered all questions, and the grouping of responses into themes reflects the Bank and the FCA's judgements about how responses can be grouped into different categories. While this FS does not refer to all of the comments in the responses (owing to their volume) nor reflect the level of detail in the responses, the Bank, the PRA, and the FCA will be making use of all the responses as we continue to consider issues around the use of AI and ML in UK financial services. Background 1.4 The supervisory authorities published DP5/22 to further their understanding and to deepen dialogue on how AI may affect their respective objectives. This is part of the supervisory authorities' wider programme of work related to AI, including the AI Public Private Forum , the final report of which was published in February 2022. 1.5 The supervisory authorities invited comments to the DP, including answers to the questions laid out in it. The DP stated that the supervisory authorities may publish a summary of the comments received. Responses 1.6 The Bank and the FCA received 54 responses from regulated firms, trade bodies, and non-financial businesses. Feedback statement structure 1.7 Chapters 2, 3, and 4 of this statement describe the responses to the corresponding chapters in DP5/22. A list of the questions posed in the DP can be found in Chapter 5.

2: Supervisory authorities' objectives and remits 2.1 Chapter 2 of DP5/22 outlined the supervisory authorities' objectives and remits, and their relevance to the use of AI in financial services. It surveyed existing approaches by regulators and authorities to distinguishing between AI and non-AI. Some approaches provide a legal definition of AI (eg the proposed EU AI Act), whereas others identify the key characteristics of AI (eg the UK AI White Paper). The DP asked respondents whether a financial services sector-specific regulatory definition is beneficial and whether there are other effective approaches that do not rely on a definition. 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? 2.2 Most respondents thought that a financial services sector-specific regulatory definition of AI would not be helpful for the safe and responsible adoption of AI in UK financial services. They offered the following reasons: (i) a definition could become quickly outdated due to the pace of technology development, (ii) definitions could be too broad (ie cover non-AI systems) or too narrow (ie do not cover all the use cases), (iii) a definition could create incentives for regulatory arbitrage, and (iv) creating a sectoral regulatory definition could be in conflict with the regulatory authorities' technology-neutral approach. 2.3 Respondents supportive of a sectoral regulatory definition suggested that such a definition may help prevent misinterpretation and/or inconsistent implementation of regulatory requirements in the use of AI in financial services. A few respondents were supportive of a broad definition but not a sector-specific one, with some noting that a sectoral definition would add to the complexity where different jurisdictions might have different definitions of AI. Some respondents suggested that whether a sectoral definition is beneficial or not would depend on the approach taken by regulatory authorities, for example, if there are AI-specific rules, then a sectoral definition would help firms understand what activities are in scope. 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? 2.4 Most respondents suggested that a technology-neutral, outcomes-based, and principles-based approach would be effective in supporting the safe and responsible adoption of AI in financial services. Respondents noted that many risks related to AI are not necessarily unique to AI itself and could therefore be mitigated within existing legislative and/or regulatory frameworks in a technology-neutral way. Respondents considered that a technology-neutral approach could align with and leverage existing approaches to financial services regulation while ensuring a proportionate balance between managing risks and supporting innovation. 2.5 With respect to a principles-based approach, respondents said that high-level principles could give firms and regulators the flexibility to adapt to technological developments. Respondents added that a principles-based approach also allows firms to tailor the identification, assessment, and management of risks to the purpose, function, and outcomes of each specific AI

use case or application. 2.6 Respondents suggested that the regulatory focus should be on the outcomes affecting consumers and markets rather than on specific technologies. They further emphasised that this outcome-focused approach is in line with the approach of existing regulation, namely, that firms should ensure good outcomes and effective oversight whether or not AI is used in the process. 2.7 Respondents also emphasised that the approach to AI should be proportionate to the risks associated with, or materiality of each specific AI application. While some respondents remarked that the existing regulatory framework should be adequate in addressing AI risks, they recognised that there may be specific risks or use cases that may require further targeted regulatory guidance or intervention.

3: Potential benefits and risks 3.1 Chapter 3 of the DP summarised the potential benefits and risks of the use of AI in financial services with respect to the supervisory authorities' regulatory objectives. It also described how the drivers of risks related to AI in financial services can occur at different levels within AI systems (ie data, models, governance etc) and how these drivers can result in a range of outcomes affecting consumer, firms, and markets. The DP invited responses on which potential risks and benefits should be prioritised (including potential risk mitigation strategies). In particular, the DP asked how AI could affect groups with protected characteristics. Q3: Which potential benefits and risks should supervisory authorities prioritise? 3.2 While respondents noted a wide range of benefits and risks, a majority cited consumer protection as an area for the supervisory authorities to prioritise. While noting that AI could provide benefits to consumers (such as driving better consumer outcomes, more personalised advice, lower costs, and better prices), respondents also said that AI could create risks such as bias, discrimination, lack of explainability, transparency, and exploiting vulnerable consumers or consumers with protected characteristics. 3.3 There was no clear consensus on the potential benefits or risks of AI for competition. While some respondents said that high barriers to entry due to resourcing (eg data and technological infrastructure) and expertise requirements could favour large incumbent firms; others argued there are no significant barriers to entry including because of the availability of open-source AI models. Some respondents noted other potential competition risks associated with AI such as tacit collusion, market manipulation, and herding behaviour. 3.4 Commenting on market integrity and financial stability, respondents highlighted that the speed and scale of AI could increase the potential for (new forms of) systemic risks, such as interconnectivity between AI systems and the potential for AI-induced firm failures. Respondents mentioned the following potential risks to financial markets: (i) emergence of new forms of market manipulation, (ii) the use of deepfakes for misinformation potentially destabilising financial markets, (iii) third-party AI models resulting in convergent models including digital collusion or herding, and (iv) AI could amplify flash crashes or automated market disruptions. 3.5 On governance, respondents suggested that the most salient risk for firms is insufficient oversight. Some respondents noted that there may not be sufficient skills and experience within firms to support the level of oversight required to ensure technical (eg data and model risks) and non-technical (eg consumer and market outcomes) risk management. Some respondents noted that a lack of technical expertise is especially worrying given the increased adoption of third-party AI software. Some respondents also pointed out the importance of human-in-the-loop for mitigating risks associated with overreliance on AI or overconfidence in the accuracy of AI. 3.6 On operational resilience and outsourcing, respondents suggested that third-party providers of AI solutions should provide evidence supporting the responsible development, independent validation, and ongoing governance of their AI products, providing firms with sufficient information to make their own risk assessment. Respondents argued that third party providers do not always provide sufficient information to enable effective governance of some of their products. Given the scope and ubiquity of third-party AI applications, respondents commented that the risks posed by third party exposure could lead to an increase in systemic risks. Some respondents said that not all firms have the necessary expertise to conduct adequate due diligence of third-party AI applications and models. Q4: How are the benefits and risks likely to change as the technology evolves? 3.7 Several respondents suggested that there is likely to be an increase in the scale of AI adoption and almost half of respondents said there is likely to be an increase in the performance and complexity of AI models in financial services. This may mean that there are benefits including increased accuracy and effectiveness of solutions, expanded scope and scale of AI applications, higher efficiency gains for organisations – all of which may lead to better consumer or market outcomes if AI is deployed in a safe and responsible manner. 3.8 Some respondents expected increased use of third-party products by regulated firms to source AI capabilities and infrastructure including AI-as-a-Service and off-the-shelf AI models. Respondents suggested that this could increase dependence on technology providers, noting that it could also increase associated risks, for example, firms may not fully understand the model (or the data it uses) and may not have sufficient controls over the risks (such as managing data bias). 3.9 Many respondents, in turn, also noted that greater scale and model complexity could lead to increased demand on governance, monitoring and oversight. Some respondents remarked that robust governance arrangements and human oversight mechanisms are required as firms become more reliant on AI

systems as inputs for decision-making or critical tasks. Some respondents emphasised the need for developing firm-wide AI skills to ensure effective AI governance. Given the pace of technological development alongside the increased complexity of models, respondents cautioned that there could be a lack of experts with sufficient understanding of the latest technologies and their risks. Q5: Are there any novel challenges specific to the use of AI within financial services that are not covered in this DP? 3.10 There was no clear consensus among respondents on the novel challenges specific to the use of AI in financial services, with some respondents remarking that there are no novel challenges at all. 3.11 Some respondents suggested that there are challenges in managing the use of open-source models and third-party AI services in the financial services sector. Respondents noted that there are difficulties in conducting sufficient due diligence and ensuring sufficient regulatory controls in using third-party AI products or services. While respondents welcomed the PRA's expectations set out in SS1/23, they pointed to areas where further clarification would be useful. For example, with wider access to off-the-shelf models, respondents suggested that there is a risk of model users not fully understanding the risks in deployment and a risk of increased reliance on third-party providers. Some respondents therefore suggested that the supervisory authorities should ensure that regulated firms provide sufficient explainability and accountability when using third-party AI models or services. 3.12 Some respondents noted that the use of AI will increase as it becomes more integrated in financial institutions' infrastructure, for example, as the technology develops, bad actors may have easier access to the tools for cyber-attacks. Respondents also expected more sophisticated cyber-attacks – attackers may develop new techniques for exploiting AI systems, hence increasing the risk of security incidents and data breaches. Through input attacks and poisoning attacks, respondents noted that attackers could disrupt an otherwise robust model. Respondents also noted that due to the opaqueness of advanced models, firms may also find it difficult to distinguish between poor model performance and a model compromised by cyber-attacks, hence reducing firms' ability to mitigate such attacks. Respondents added that if this becomes widespread, it may in turn create systemic risks to financial markets. 3.13 Respondents suggested that as the technology develops, there may also be increased access of AI tools by bad actors who wish to use AI for fraud and money laundering. For example, respondents noted that generative AI can easily be exploited to create deepfakes as a way to commit fraud. The technology may make such fraud more sophisticated, greater in scale and harder to detect. This may in turn create risks to consumers and, if sufficient in magnitude, financial stability. 3.14 Some respondents noted that the adoption of Generative AI (GenAI) may increase rapidly in financial services. Respondents noted that the risks associated with the use of GenAI are not fully understood, especially risks related to bias, accuracy, reliability, and explainability. Due to 'hallucinations' in GenAI outputs, respondents also suggested that there may be risks to firms and consumers relying on or trusting GenAI as a source of financial advice or information. Q6: 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? 3.15 Most respondents suggested that the risk of bias (intended or unintended), discrimination, and financial exclusion may be particularly relevant for consumers with protected characteristics or characteristics of vulnerability. Most respondents considered the consumer harms associated with AI originate mostly from the data. Respondents further suggested that data bias, and unavailability of sufficient key data are drivers of consumer harms. To ensure good consumer outcomes, most respondents agreed that data used to build an AI system should be representative, diverse, and free from bias. 3.16 Some respondents noted that AI could help mitigate consumer harms. For example, AI has the potential to help firms to identify patterns of unfair or discriminatory outcomes more effectively, to create new products or services that cater to consumers with characteristics of vulnerability, and to enhance financial inclusion. 3.17 On mitigating consumer harms, most respondents suggested that firms should focus on mitigating data bias. Respondents suggested mitigation strategies such as addressing data quality issues, documenting biases in data, and capturing additional data that may highlight the impact on particular groups with shared characteristics. Some respondents suggested firms monitor and evaluate consumer outcomes if an AI model is used to deliver regulated products or services. Some respondents emphasised the importance of effective governance and oversight surrounding data and AI in ensuring good consumer outcomes, especially in preventing bias and discrimination. 3.18 On how the supervisory authorities could mitigate consumer impact, respondents suggested that the supervisory authorities could release guidance to clarify regulatory expectations. Some respondents welcomed further guidance on the interpretation and evaluation of good consumer outcomes in the AI context with respect to existing sectoral regulations such as the FCA's Consumer Duty. Some respondents suggested guidance on preventing, evaluating, and mitigating bias, with case studies to help illustrate best practice. Some respondents suggested guidance on the use of personal data in AI in the financial services context, supported by case studies to demonstrate what good looks like. Q7: 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? 3.19 A number of respondents stressed that the most relevant metrics

for assessing the benefits and risks of AI in financial services would depend on the specific use case. For example, the benefits of an AI model used to improve payment matching could be assessed using metrics such as: the percentage increase in the number of payments processed; the percentage reduction in processing errors; and improvement in customer satisfaction. By contrast, in anti-money laundering applications, two key performance metrics of an AI-based system might be (i) precision (the proportion of highlighted records that are ultimately found to be suspicious), and (ii) recall (the proportion of suspicious records that are highlighted by the system). One respondent also noted that relevant metrics would depend on whether the AI system is processing personal, commercial, market or other data. In addition, one respondent noted that firms should explore multiple, complementary metrics for every application of AI, to ensure they have a comprehensive view of the benefits and risks. A few respondents argued that it was too early to say which metrics would be most relevant, underscoring the need for continued engagement between regulators, practitioners, and academia. 3.20 Notwithstanding the specifics of the intended use case, around a half of respondents to the question noted that metrics focused on consumer outcomes would be most important in assessing the benefits and risks of the use of AI. One respondent noted that indicators of better outcomes for customers might include the factors already set out in the FCA's Consumer Duty eg evidence of more consumer-centric product or service design, evidence of engagement with customers so they can make effective, timely and properly informed decisions about financial products and services, and evidence that consistently considering the needs of their customers, and how they behave, at every stage of the product/service lifecycle. Others suggested indicators that included customer satisfaction and complaint metrics. Some respondents in turn noted that outcome-based metrics should not differ across AI and non-AI systems. 3.21 Also related to consumer outcomes, several respondents noted the importance of monitoring fairness metrics ie those designed to identify biased outcomes. A number of respondents noted that there are quantitative metrics to evaluate the fairness of an AI system's outcomes, including: equalised odds, demographic parity, disparate impact, statistical parity difference, average odds difference, and equalised opportunity difference. However, some respondents also noted that each of these metrics has a different focus, and ultimately rely on a different concept of fairness. One respondent argued that regulatory authorities therefore needed to engage with industry to establish which metrics are most appropriate within each context. 3.22 Aside from metrics related to consumer outcomes, around a half of respondents stressed the need for data and model performance metrics, both during the development of the model and after deployment, to build a comprehensive view of the risks and benefits. Suggested areas where metrics would be important included data quality, data representativeness, data drift, model accuracy, model robustness, model reproducibility, model/concept drift, the time and cost required to train and re-train the model. Several responses noted the need for firms to use explainability metrics, for example Shapley Additive Explanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME). One respondent flagged the need to measure the extent of the system's autonomy from human control/oversight, and the ease with which a system could be turned off without causing disruption to business services.

4: Regulation 4.1 Chapter 4 of the DP discussed the parts of the current regulatory framework for financial services that are most relevant to the regulation of AI. The chapter explained the supervisory authorities' approach to the risks identified in Chapter 3. The DP invited feedback on whether clarification of existing regulatory requirements and expectations would be helpful, and how the supervisory authorities could best support the safe and responsible adoption of AI in UK financial services. Q8: Are there any other legal requirements or guidance that you consider to be relevant to AI? 4.2 A number of respondents pointed to other legal requirements or guidance. One respondent noted that the suite of regulation governing the use of AI in the financial sector is already extensive, so care is needed to avoid creating unnecessary new requirements. Another highlighted the importance of a consistent supervisory approach to the use of AI across the regulatory authorities. 4.3 Respondents highlighted legal requirements and guidance related to data protection. One response noted regulatory guidance indicating that the 'right to erasure' under the UK General Data Protection Regulation (UK GDPR) extends to personal data used to train AI models, which could prove challenging in practice given the limited extent to which developers are able to separate and remove training data from a trained AI model. One respondent argued that, although it is generally recognised that data protection laws apply to the use of AI, there may be a lack of understanding by suppliers, developers, and users, leading to those actors potentially gaming or ignoring the rules. 4.4 A number of respondents stressed the relevance and importance of the existing regulatory framework related to operational resilience and outsourcing, including the PRA's supervisory statements (SS) 1/21 – Operational resilience: Impact tolerances for important business services and SS2/21 – Outsourcing and third party risk management , as well as the FCA's PS21/3 – Building operational resilience . Respondents also noted the relevance of the Bank, the PRA and the FCA's DP3/22 – Operational resilience: Critical third parties to the UK financial sector . 4.5 Respondents argued that an industry-wide standard for data quality is necessary and a

potential legal gap, as current legal requirements, and regulatory expectations applicable to regulated firms do not include a common standard for data quality. 4.6 Other laws and guidance considered relevant by respondents included: discrimination laws (eg the Equality Act 2010), intellectual property law, contract law, laws relating to the use of electronic communications and infrastructure (eg the Product Security and Telecommunications Infrastructure Act 2022), and forms of ethical guidance. 4.7 Some respondents noted legal and regulatory developments in other jurisdictions (including the proposed EU AI Act), and argued that international regulatory harmonisation would be beneficial, where possible, particularly for multinational firms. One respondent noted that the development of adequate and flexible cooperation mechanisms supporting information-sharing (or lessons learnt) across jurisdictions could also minimise barriers and facilitate beneficial innovation. 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? 4.8 Many respondents thought that there were no regulatory barriers to the safe and reasonable adoption of AI in UK financial services. 4.9 A number of respondents disagreed and referred to regulatory barriers around data protection and privacy. Some respondents addressed the interaction between AI and the UK GDPR. One respondent noted that the way UK GDPR interacts with AI was hard to navigate and referenced the UK GDPR provisions on rights related to automated decision making including profiling (Rights related to automated decision making including profiling). Another viewed the interaction between AI and the UK GDPR as amounting to a prohibition on automated decision-making. A few respondents said that data localisation requirements were a concern. For example, they noted that the EU's proposed Cybersecurity Certification Scheme on Cloud Services would include mandatory localisation in the EU. Q10: How could current regulation be clarified with respect to AI? 4.10 Several respondents sought clarification on what bias and fairness could mean in the context of AI models, more specifically, they asked how firms should interpret the Equality Act 2010 and the FCA Consumer Duty in this context. Other respondents asked for more clarity on how data protection / privacy rights interact with AI techniques, echoing the responses to Q9. Another group of respondents asked for more clarity on explainability for AI applications. 4.11 Several respondents asked for more clarity on the issue of governance and human oversight and where accountability for the outcome of AI applications sits within a firm. Some respondents linked this to the need for further regulatory clarity on the use of third-party vendor solutions. Another set of respondents requested clarity on how firms should assess and rate AI model risk for the purposes of model risk management. One respondent argued that the tiering for ratings should be aligned with the definition of 'important business service' as defined in the Operational Resilience Part of the PRA Rulebook, and that firms should maintain an inventory of high-risk AI use cases. Q11: How could current regulation be simplified, strengthened and/or extended to better encompass AI and address potential risks and harms? 4.12 Some respondents asked for more practical and actionable guidance and illustrative case-studies. Others called for more direct regulatory guidance on risk-based approaches to AI ie that a firm's risk management measures (including data, model risk management, and governance) should be proportionate to the risk posed by the application/use-cases of AI techniques. 4.13 A few respondents requested more clarity on implementing bias/fairness requirements in practice. Others requested more clarity on the usage of third-party vendors, and some called for a consolidated stocktake of relevant legislation and regulation to better understand how existing legislation and regulation encompass AI. 4.14 Some respondents called for clarity on requirements relating to: accountability and data protection. Two respondents noted the importance of regulatory sandboxes in facilitating responsible innovation. Q12: Are existing firm governance structures sufficient to encompass AI, and if not, how could they be changed or adapted? 4.15 Most respondents said that existing firm governance structures are either already sufficient to cover AI or are being adapting by firms to make them sufficient and to comply with existing regulatory requirements. 4.16 Some respondents said that firms should implement a central or strategic AI function or committee to approve and/or oversee AI deployment across the firm and ensure a coherent approach. Others argued that local business areas should retain accountability of specific AI applications and outputs although notably this was not seen as a mutually exclusive option – a firm might have both a central function and accountability within specific business areas. Some respondents noted that board-level or senior management expertise or oversight of AI adoption is necessary. A few respondents said that governance should be treated proportionately (eg smaller firms may struggle to put in place the same structures as larger firms). Finally, a few respondents thought it is important to embed relevant data science skills in audit and compliance functions. 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? 4.17 Most respondents did not think that creating a new Prescribed Responsibility (PR) for AI to be allocated to a Senior Management Function (SMF) would be helpful for enhancing effective governance of AI. Several respondents argued that there were too many potential applications of AI within a firm and/or that the relevant responsibilities were already reflected in the PRs or could

be reflected in the 'statements of responsibilities' for existing SMFs. Many respondents also said that firms need local owners to retain accountability over AI, while some said that this approach is not technology neutral. Two respondents stated that adding this PR could overburden the Chief Operations Officer (SMF24), that being the most likely SMF to be allocated a PR for AI. 4.18 Other respondents disagreed and believed that it would be helpful to create a new PR for AI to be allocated to an SMF. Respondents argued this would create an incentive for meaningful accountability for AI deployment and oversight within firms. Two responses suggested that such a PR could make the SMF a helpful single point of contact for regulators. Q14: Would further guidance on how to interpret the 'reasonable steps' element of the SM&CR in an AI context be helpful? 4.19 Most respondents thought that further guidance on how to interpret the 'reasonable steps' element of the SM&CR in an AI context would be helpful, although only if it was practical or actionable guidance. Three respondents stressed the importance of engaging with industry on the substance of that guidance. 4.20 Some respondents disagreed, stating that further guidance would not be helpful. Three noted that additional guidance could create duplication or confusion with respect to existing regulation and guidance. Q15: 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? 4.21 Most respondents argued that there are areas of data regulation that are not sufficient to identify, manage, monitor, and control the risks associated with AI models. Some pointed to insufficient regulation on the topics of data access, data protection, and data privacy (eg to monitor bias). Some respondents thought that regulation in relation to data quality, data management, and operations are insufficient. 4.22 Many respondents said that there would be value in alignment of the supervisory authorities' data definitions / taxonomies. This theme was consistent with a more general call for greater coordination and harmonisation among sectoral regulators. Several respondents noted that data regulation is a patchwork, and that there should be a more consolidated approach. Two respondents felt that data regulation is already sufficient. Q16: 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? 4.23 Some respondents suggested that the supervisory authorities clarify regulatory expectations through additional guidance or guidance on best practice. Respondents remarked that, while existing regulation is sufficient to cover risks associated with AI, there are areas where clarificatory guidance on the application of existing regulation is needed (such as accountability of different parties in outsourcing) and areas of novel risk that may require further guidance in the future. Some respondents suggested that guidance on best practices for responsible AI development and deployment would help firms ensure that they are adopting AI in a safe and responsible manner. 4.24 Several respondents suggested that the supervisory authorities focus on consumer protection in developing their regulatory approach. They encouraged the supervisory authorities to ensure firms act with integrity and in consumers' best interests by promoting principles such as fairness, transparency, and explainability. Respondents emphasised that managing data bias and having effective oversight over the development process is key to safe and responsible deployment of AI, especially for ensuring good outcomes for vulnerable consumers. Some respondents suggested that the supervisory authorities should explore ways to identify and address bias, discrimination, and exploitation, including monitoring and testing outcomes for consumers. 4.25 Many respondents emphasised the importance of cross-sectoral and cross-jurisdictional coordination as AI is a cross-cutting technology extending across sectoral boundaries. Respondents noted that certain AI use cases are cross-sectoral and that there are regulatory and legislative provisions beyond financial services relevant in the adoption of AI, such as those for data protection and equality. 4.26 Therefore, respondents encouraged authorities to ensure coherence and consistency in regulatory approaches across sectoral regulators, such as aligning key principles, metrics, and interpretation of key concepts. Some respondents suggested that the supervisory authorities work with other regulators to reduce and/or prevent regulatory overlaps and clarify the role of sectoral regulations and legislation. Respondents added that since many regulated firms operate in multiple jurisdictions, an internationally coordinated and harmonised regulatory response is critical in ensuring that UK regulation does not disadvantage UK firms and markets while also minimising fragmentation and operational complexity. 4.27 Many respondents suggested that the supervisory authorities collaborate and/or set up working groups with industry, academia, and civil society. Promoting transparent and collaborative conversations with these stakeholders, respondents suggested, would help address and monitor issues that cut across the AI lifecycle. Respondents noted that further ongoing public-private engagement would also help the supervisory authorities to ensure that regulatory approaches adapt to changing technological developments and public expectations. Q17: Which existing industry standards (if any) are useful when developing, deploying, and/or using AI? Could any particular standards support the safe and responsible adoption of AI in UK financial services? 4.28 Respondents suggested a range of international and other standards in response to this question. On international standards, respondents raised examples of such as BCBS 239

Principles for effective risk data aggregation and risk reporting , BCBS 328 Corporate governance principles for banks , US National Institute of Standards and Technology's AI Risk Management Framework . 4.29 Some respondents remarked that emerging international standards would be important in establishing a coherent global approach in promoting AI innovation while proportionately managing risks, and ensuring global interoperability. Some respondents cautioned that, given the rapid pace of AI development, introducing additional standards at this stage could create uncertainty and in turn may stifle innovation. 4.30 The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) have jointly formed a committee to develop international standards on Artificial intelligence, ISO/IEC JTC 1/SC 42 – Artificial intelligence . Respondents cited the ongoing developments of standards by this committee to support industry understanding, including ISO/IEC TS 4213 Assessment of machine learning classification performance , ISO/IEC TR 24027:2021 Bias in AI systems and AI aided decision making , ISO/IEC 24029 series Assessment of robustness of neural networks . 4.31 In addition, there are other mentions of related international technical standards for example on risk, security: ISO/IEC 27001 Information security management systems, ISO/IEC 15504 series (all parts) provides a framework for assessment of processes in systems and software engineering, ISO 31000 Risk management, ISO/IEC 27032 Cybersecurity - Guidelines for Internet security, ISO/IEC/IEEE 29148 Systems and software engineering — Life cycle processes — Requirements engineering. The Institute of Electrical and Electronics Engineers (IEEE) has also developed standards for Artificial Intelligence (AI) Model Representation, Compression, Distribution, and Management. 4.32 Some respondents pointed to the AI Standards Hub, a partnership initiative established jointly between the Alan Turing Institute, the British Standards Institution, the National Physical Laboratory, and HM Government. 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? 4.33 A number of respondents argued that a risk-based and principles-based regulatory approach to the use of AI in financial services is key, and would better support international interoperability than the adoption of more prescriptive – and potentially conflicting – rules. 4.34 In that vein, several respondents said that there was value in the draft EU AI Act's risk-based categorisation of AI use-cases (eg unacceptable risk, high risk, low or minimal risk), whereby regulatory requirements were proportionate to the risks posed by the particular use case and the potential for harm. One respondent thought this would enable a consistent classification of use-cases across firms and help firms focus their risk management on higher risk applications. On the other hand, another respondent said that the EU's draft proposal focuses on a single risk criterion (impact on humans), which might not be the appropriate focus for prudential regulators. The respondent gave the example of algorithmic trading, which may not be classified as high risk under the EU's draft proposals but could be seen as high risk from a prudential or financial stability perspective. Another respondent also noted that any future framework must be aligned and work in conjunction with other sector-specific regimes that currently provide for the regulation of digital technologies or will do so in the future. They noted an example highlighted by MedTech Europe of certain duplications between the proposed EU AI Act and obligations within the EU Medical Device Regulations and In Vitro Device Regulations . 4.35 Several respondents also pointed to risk-based guidance from the US. A number of respondents noted that the US Federal Reserve Supervisory Letter SR 11-7 on guidance on Model Risk Management is an effective framework for managing model and AI risk. 4.36 Two respondents also highlighted the National Institute of Science and Technology's AI Risk Management Framework: one respondent said that this is a helpful approach, both for setting out key characteristics of trustworthy AI systems, and providing a playbook of actions for achieving the specified outcomes. 4.37 Several respondents pointed to the Monetary Authority of Singapore's (MAS) FEAT Principles (Fairness, Ethics, Accountability and Transparency) and Veritas consortium programme with industry. One respondent said that these initiatives provide a potential path forward in developing and operationalising principles for the use of AI in financial services. Another respondent argued that the MAS's approach balances governance requirements with practical considerations. One respondent noted that the MAS had published a number of use-case studies demonstrating how certain organisations have successfully integrated the FEAT Principles and the initial roadblocks met along the way. The respondent also highlighted the open-source software toolkit that MAS has released to supplement its publications, adding that it could alleviate early technology challenges faced by many – particularly smaller – organisations when looking to use AI. 4.38 A number of respondents stressed the importance of harmonising regulatory approaches to AI across borders to mitigate the risk of conflicting requirements and significant compliance costs for cross-border firms and those wishing to expand into new markets. Q19: Are there any specific elements or approaches to apply or avoid to facilitate effective competition in the UK financial services sector? 4.39 Many respondents suggested that addressing regulatory uncertainty could help facilitate effective competition. Clear regulatory expectations and guidance could give firms confidence to innovate, especially for smaller firms which tend to struggle disproportionately from regulatory uncertainty. Many respondents suggested that regulatory approaches could

help facilitate competition by fostering a level-playing field for firms to innovate. Such approaches would be sufficiently flexible to adapt to technological change and manage public expectations. Respondents added, that overly prescriptive or restrictive regulations would, on the other hand, hinder the development of AI and are likely to become ineffective due to the pace of technological progress. 4.40 Several respondents also emphasised that the regulatory approach should be proportionate and not impose excessive costs on firms. High costs and regulatory burdens could be a barrier to entry and could potentially widen the gap in innovation between regulated and non-regulated financial services firms. 4.41 Many respondents suggested that the supervisory authorities should support collaboration between financial services firms, regulators, academia, and technology practitioners with the aim of promoting competition. Respondents also noted that encouraging firms to collaborate in the development and deployment of AI, such as sharing knowledge and resources, could help reduce costs and improve the quality of AI systems for financial services. This could also help make innovation more available to smaller firms where they may lack resources or capabilities. Engaging with a wide range of stakeholders would help ensure that regulatory proposals and frameworks are adaptable and proportionate to different types of businesses in financial services. 4.42 Some respondents suggested that open banking could help improve data access within financial services and thus facilitate innovation with AI and competition. Lack of access to high-quality data may be a barrier to entry for firms' adoption of AI. Open banking may help create a more level playing field by providing firms with larger and more diverse datasets, and therefore enabling more effective competition.

5: Questions The questions listed below are all the questions that appeared in DP5/22. 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? 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? Q3: Which potential benefits and risks should supervisory authorities prioritise? Q4: How are the benefits and risks likely to change as the technology evolves? Q5: Are there any novel challenges specific to the use of AI within financial services that are not covered in this DP? Q6: 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? Q7: 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? Q8: Are there any other legal requirements or guidance that you consider to be relevant to AI? 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? Q10: How could current regulation be clarified with respect to AI? Q11: How could current regulation be simplified, strengthened and/or extended to better encompass AI and address potential risks and harms? Q12: Are existing firm governance structures sufficient to encompass AI, and if not, how could they be changed or adapted? 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? Q14: Would further guidance on how to interpret the 'reasonable steps' element of the SM&CR in an AI context be helpful? Q15: 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? Q16: 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? Q17: Which existing industry standards (if any) are useful when developing, deploying, and/or using AI? Could any particular standards support the safe and responsible adoption of AI in UK financial services? 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? Q19: Are there any specific elements or approaches to apply or avoid to facilitate effective competition in the UK financial services sector?

Convert this page to PDF

Other prudential regulation releases

Prudential Regulation // Policy statement

14 October 2026

PS19/26 – High loan to income lending

PS19/26 – High loan to income lending

Prudential Regulation // Supervisory statement

14 October 2026

SS2/26 – High loan to income lending

SS2/26 – High loan to income lending

Prudential Regulation // Consultation paper

07 October 2026

CP13/26 – The Prudential Regulation Authority's...

CP13/26 – The Prudential Regulation Authority's automatic thresholds indexation framework

Prudential Regulation // PRA Regulatory Digest

01 October 2026

PRA Regulatory Digest – September 2026

PRA Regulatory Digest – September 2026

[View more Other prudential regulation releases](#)

[Back to top](#)