

UK Finance response to DP5/22 – Artificial Intelligence and Machine Learning

Date: 10 February 2023

Sent to: DP5_22@bankofengland.co.uk

UK Finance is the collective voice for the banking and finance industry.

Representing more than 300 firms across the industry, we act to enhance competitiveness, support customers and facilitate innovation.

The responsible uptake of artificial intelligence (AI) by the financial sector has the potential to provide great benefits for consumers and businesses in the UK. Nonetheless, we recognise that these technologies do pose certain risks, which industry and regulators are rightly working to understand and address.

Artificial intelligence brings great opportunities for consumers, businesses and the UK economy, many of which are mentioned in discussion paper DP5/22 ('the paper'). We recognise that AI also brings with it certain risks, which regulators and firms need to be alive to. The paper is an important step towards an effective and proportionate regulatory approach, building on the valuable work of the AI Public-Private Forum.

Please find annexed our responses to the paper's questions.

Summary of our key points:

- The best approach to financial sector AI regulation is not to have an 'AI overlay' of new rules, hinging off a regulatory definition of AI. This would create a high risk of duplication and added complexity. Similarly, we see an important role for voluntary standards but would not consider mandatory standards to be an effective regulatory approach.
- Instead, we should maintain a technology-neutral, outcomes-focused and principles-based regulatory framework. Where novel risks emerge over time, or where existing risks are amplified by the use of AI, guidance targeting the specific types of system and use case in question will most often be the appropriate regulatory response.
- We see the existing regulatory framework as being broadly fit for purpose already. We do not think a new Senior Management Function for AI is appropriate, given the different structures firms will have.
- Overall, we do not see the current regulatory framework as creating significant barriers to responsible AI adoption in the UK, although there are some areas where guidance may be beneficial in due course once the overall regulatory framework has become clearer. That said, we do see a risk of new rules multiplying, leading to confusion and duplication, particularly in situations where multiple regulatory authorities are competent.
- We emphasise the importance of coordination with other authorities. Many key AI risks are not unique to financial services, vendors may offer products into multiple markets and cross-sectoral use cases may emerge. Regulatory cooperation will help mitigate risks of regulatory arbitrage, duplication of rules and unwarranted differences of approach between sectors.

If you have any questions relating to this response, please contact:

- Walter McCahon, Principal, Privacy and Data Ethics
walter.mccahon@ukfinance.org.uk
- Simon Hills, Director, Prudential Policy
simon.hills@ukfinance.org.uk
- Harriet Wilson, Analyst, Digital Technology and Cyber
harriet.wilson@ukfinance.org.uk

Walter McCahon

Principal, Privacy and Data Ethics

Annex: Responses to discussion paper questions

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

This depends on the approach to AI taken by authorities.

If, ultimately, there is to be an 'AI rulebook' for financial services, then a definition will likely be required so that firms understand what activities or technologies are in scope.

However, as we expand on in other sections of our response, the 'AI rulebook' option may not be the best approach to managing AI risks. It has proven difficult to agree a satisfactory definition of AI for regulatory purposes where this has been attempted, in part because some 'AI risks' are not unique to AI and / or are not relevant to all systems that are sometimes termed 'AI'. As such, it would be difficult to create a definition that simultaneously:

- does not have gaps,
- avoids catching simple systems that do not create novel risks (or materially amplify existing risks), and
- does not create duplication.

Therefore, tying regulation to a single 'AI' definition may detract from the risk-based approach.

Such an approach would bring further challenges if there were multiple global, UK and cross-sectoral definitions. Having multiple varying definitions would create unnecessary complexity, tensions and conflicts, especially where use cases cross jurisdictions or sectors.

We also note that if a definition were constructed, it would need to allow for the evolution of AI over time, putting it at risk of going out of date.

Therefore, instead of an 'AI overlay' we recommend an approach that targets novel or amplified risks in a more focused manner, as outlined below under question 2.

That said, *if* authorities do eventually decide that a regulatory AI definition is needed, definitions could be looked at from among those contained within regulatory and industry papers that generally focus on a technique/algorithm-based definition. The example put forward by the [OECD](#) to explain an 'AI system' is one prominent option that could be used as a baseline, and would/could minimise tensions and overlaps between countries and sectors.

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?

In our view, a set of 'AI regulations', or an 'AI overlay', is not the most effective way to address risks associated with the use of such technologies in the financial sector. The focus of regulation should be on outcomes and specific risks, making use of high-level principles. Such a technology-neutral approach should best ensure that risks and customers' needs are addressed effectively, irrespective of the technical means of service delivery. For example, firms should ensure fair (unbiased) outcomes for consumers and ensure effective oversight of systems, whether or not AI is involved.

This is already the basis of the existing regulatory frameworks, which we see as being effective for AI governance.

We recognise that existing rules and guidance may not address new 'AI issues' that may arise over time but – on the whole – AI tends to exacerbate pre-existing risks, rather than create new ones. As such, it makes sense to consider AI risk within the context of existing risk categories. Making tactical additions or amendments to these existing frameworks is likely to be sufficient to resolve emerging issues, rather than requiring an 'AI overlay'. In most cases, it is likely that guidance on existing rules, rather than the creation of new rules, will be an appropriate regulatory response to emerging issues. Guidance also has the benefit of being more flexible and adaptable over time than hard rules.

In this context, the fresh guidance would have to specify the nature of the risks it is intended to mitigate, which might not be relevant to all 'AI technology' or indeed all use cases. For example, guidance relating to systems that exhibit autonomy and adaptability could be needed in relation to certain governance rules in due course, as might guidance on the oversight of opaque systems. However, these should be focussed on particular areas or applications giving rise to such new specific risks and would not be relevant to all techniques that can be classified as 'AI'.

We note, of course, that such guidance would need to be clear in its scope, likely requiring the definition of key terms as the need arises. Terms like 'autonomous', 'adaptable' and 'opaque' could be examples, with potential to align with the central work on the overarching [AI regulatory approach](#) being led by the Office for AI, which is provisionally focusing on the risks stemming from autonomy and adaptability of systems.

In our view, this more targeted approach would be more risk-based than an approach relying on new rules or guidance that apply to a broad definition of 'AI', better focusing firms' and authorities' attention on the most important issues.

Benefits, risks, and harms of AI

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

Risks

Noting, first, that saying certain issues should be prioritised does not necessarily mean that fresh regulation is required, we think that some of the key areas of risk that should be prioritised by regulators are:

- Unfairness, discrimination and the addition / removal of bias.
- Explainability and the 'black box' challenge.
- The potential for 'human agency' to be undermined, either through customer manipulation or through automated decisions that are not explained and cannot be appealed against (connects to explainability).
- The potential for the scale permitted by AI technologies to amplify the extent of issues within a single firm, or to create systemic risks, such as competition and collusion risks.
- Risks relating to the use of untested or novel data sets (noting that this goes beyond AI).
- Outsourcing issues, such as IP and other commercial challenges that can arise, including the potential lack of visibility for firms wishing to use an external AI provider.

These risks should be looked at in the context of specific areas of AI application, as the implications of – for example – bias or explainability issues will vary depending on the context.

Existing Regulation:

Whilst these are key risk areas, existing regulatory frameworks do already cover them so there is a risk of duplication and overlap if further rules are added in future. In particular we note:

- The updated model risk framework from the PRA.
- The Principles for Business and Consumer Duty, which provide overarching obligations on firms to consider communications and fairness, including the fairness of any 'nudges'.
- Any regulatory financial sector guidance or expectations in relation to discrimination would need to dovetail with the perspective of the Equality and Human Rights Commission, while guidance on explainability and automated decision-making would need to dovetail with UK General Data Protection Regulation (GDPR) and the views of the Information Commissioner's Office (ICO), insofar as personal data is involved.
- With regards to the risk above on systemic risks connected to competition and collusion, although this is an important area to monitor, we note that there are already rules in place for algorithmic trading, which were tightened following the 'mini crash' of 2015.

We also note that the Bank of England (BoE) and Financial Conduct Authority (FCA) would need to coordinate with other authorities (perhaps through the Digital Regulation Cooperation Forum) on most – if not all – of the key risks set out above.

In considering risks and ethics, it is important to consider all three elements identified in the discussion paper: the data, the model and the governance. Sometimes discussions and thought pieces can focus too narrowly, for example on 'algorithmic fairness', which can risk losing sight of the importance of fairness and ethics questions associated with the data.

Benefits

Whilst there are risks that need to be prioritised, consideration must also be had to the benefits that AI may bring, when determining the regulatory approach.

Some key benefits of AI include:

- Improvements to decision quality in terms of consistency and relevance.
- Improvements to response time, which improves client outcomes.
- Potential to improve business compliance with legal and regulatory requirements, if implemented appropriately.
- Potential to anticipate issues and mitigate risks.

AI technologies also bring the potential for *more* explainability and *less* unfairly biased decision-making than where decisions are taken by human beings. As AI techniques develop, new types of system may be able to adjust themselves to counter issues, such as by auto-detecting bias. The flip side of these key AI risks is that, if they are well managed and overseen, decision-making can be made more transparent and fairer for consumers. This offers the potential to improve financial inclusion.

We expand further on benefits under the next question.

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

As the technology evolves, AI systems are likely to become more autonomous and data sets are likely to become bigger and more varied. Therefore, there is likely to be an amplification of both risks and benefits.

Risks

Risks surrounding AI will likely 'scale up', with the potential impact becoming broader as use cases expand. It is also possible that market collusion and systemic events will become more likely. That said, and as noted above, existing rules do already sit over the top of these risks.

There will be more complexity in oversight, monitoring and governance as use cases become more diverse, more complex models are used, and more chains of AI are built with AI outputs from one system fed into other systems. However, due to this 'scale up', it may also become easier to identify emerging issues.

As wider, more novel, data sources become available over time, we may see an increase in the salience of questions around how to ensure fairness when certain types of data start being used in new contexts or to provide new types of services. This connects to the risks identified in the paper connected to privacy, risk pooling and the potential for financial exclusion. However, GDPR does already provide relevant safeguards against these risks, notably its rules on purpose limitation, compatibility, fairness, transparency and 'legal basis', so coordination with the ICO would be needed in relation to these issues.

Benefits

As AI develops, and improves, the benefits of AI should also evolve, ultimately leading to better outcomes for businesses and consumers. For example:

- It is likely that response times and the resolution of problems will become quicker and more efficient, as expertise and familiarity with techniques develop.
- An increasing number of repetitive and time-consuming tasks will be automated using intelligent solutions, increasing efficiency and reducing costs
- The increasing ability to make use of larger datasets and to better connect and integrate different datasets within financial services firms is likely to improve the quality of decisions made with AI.
- There will be increased integration of AI into businesses' compliance with legal and regulatory requirements.
- Increasing sophistication of AI is likely to enable more powerful tools to scrutinise or track the operation and deployment of AI solutions, allowing greater control over the risks highlighted above.

We also note several broader trends that will impact likely benefits:

- In addition to the implications for *risks* noted above, we also highlight that increasing scale and scalability may also bring *benefits* over time. As technology evolves and use of AI becomes more widespread, it could provide firms with economies of scale which enable them to offer their products and services to a wider group of customers within their target market, in a cost-effective manner which can translate into better value of the customer. For example, making credit decisions more efficient and cheaper for both the firm and the customer.
- As more use cases emerge, this could enable more opportunities for more players in the market, increasing competition and giving consumers more choice.
- Further development of open banking and, later, smart data will presumably enable more AI use cases in the area of switching, personalisation, improved product / price finding and advice, boosting the competition benefits of AI. However, this also brings privacy and data security risks.
- The market and authorities may gain confidence in the use of more opaque systems in the future, provided that these are appropriately tested and monitored, as seen in other regulated industries, such as pharmaceuticals. This would over time enable firms and consumers to gain more from the benefits offered by more opaque

systems, such as accuracy. Further thought on how to develop this greater confidence will be needed.

Overall, these benefits will enhance decision-making and present more choice in the market, leading to better outcomes for both businesses and consumers.

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

The discussion paper covers the main risks well.

One additional challenge to highlight relates to regulators' expectations for firms to use AI to achieve public policy or regulatory objectives. Where regulators create rules or guidance that would have the effect of obliging firms to make greater use of AI, they should consider any trade-offs that this might create. In particular, we could see in the future growing expectations for AI to be used by firms to better understand consumers' individual needs and vulnerabilities, etc, or to enable profiling of customers to identify risks associated with specific cohorts or demographic groups. However, these uses can be intrusive, bringing an ethical / privacy trade-off, particularly if they serve to reduce choice (eg: if various parties use AI and the AI comes to the same conclusions as to what is the best course of action or product / service offering).

Regulators should tackle these trade-offs directly and publish their thinking when considering any such future regulatory developments, eg, further development of the Consumer Duty or vulnerable customer guidance. They should engage with consumers to test their preferences and views of these trade-offs. Collaboration with the ICO and Centre for Data Ethics and Innovation (CDEI) would assist in making these assessments. Considering these trade-offs will hopefully improve regulatory clarity in these areas and reduce the likelihood of firms needing to navigate conflicting regulatory positions.

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?

The risk of AI systems causing discrimination is well documented.

There are a number of actions that firms ought to take in order to mitigate this risk effectively, including.

- Ensuring appropriate and representative training data is used.
- Validation and testing, for example using dummy data and then reviewing outputs.
- Considering what variables are appropriate to use.
- Considering potential differential impacts on different protected characteristic groups in product and process design.
- Monitoring of live systems.
- Use of appropriate fairness metrics.

There are some challenges firms face in taking these actions, however:

- There is uncertainty over the appropriate choice of fairness metrics. Many of these have been developed in academic settings over recent years but there is no definitive means to determine which is appropriate in any given situation. Firms therefore need to make a reasoned decision based on the use case and outcomes for customers, likely including, for example, the relative impacts on customers of false positives vs false negatives. This choice needs to be within the legally binding guard rails of the

Equality Act framework, notably including the prohibition on 'positive discrimination', but recognition of 'positive action'.

- (We further note that this challenge with choosing a fairness metric and interpreting the tests under the Equality Act is not unique to AI. Many metrics have emerged via the AI fairness debate, but they could apply more widely. Presumably if – hypothetically – a certain difference in positive rates for demographic group 1 vs demographic group 2 is unfair or discriminatory, this would be true whether the decisions are made by an AI system, a simpler decisioning systems or indeed stem from purely human decision-making).
- In addition, efforts to measure differences in outcomes between different demographic groups are limited by data protection rules for 'special category data' collection, combined with likely customer reluctance to disclose demographic data to firms. There is potential for synthetic data becoming available for testing and validation which might assist with both challenges, however. DCMS has also signalled that UK data protection rules could be amended to facilitate model validation and testing fairness (though this may not reduce customer reluctance).

Subject to the 'special category data' constraint noted above, AI also offers the potential to enable firms to identify potential unfair or discriminatory patterns of outcomes more effectively. There is therefore a potential for AI uptake to *reduce* unfairness and discrimination.

With regards to regulators' actions:

- Despite the uncertainty over fairness metrics described above, we do not think that this requires a response from BoE and FCA at this time. We note that the Equality and Human Rights Commission has already flagged in their three-year strategy that they will work on AI discrimination guidance, so any financial sector guidance would need to align with its views.
- In any event, in line with our wider comments in our response about the regulatory approach to AI, guidance would need to focus on outcomes and principles and would need to avoid prescriptively determining metrics. Which metric is most appropriate will depend on a range of factors, such as the use case and the target market.
- With regards to controls, we similarly do not see a gap to be filled at this time, as firms already need to have controls in place to ensure compliance with existing rules, such as the Principles for Business and the incoming Consumer Duty, as discuss in the discussion paper in 4.16 and elsewhere.

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?

Appropriate metrics may be quite diverse, depending on the type of use case and whether the AI system is being used to make or inform decisions about individuals, about institutional clients or about internal matters that do not directly touch customers. Similarly, the relevant metrics will depend on whether the system is processing personal, commercial, market or other data.

Specifically in relation to AI applications that impact individuals, building on comments above in relation to protected characteristics groups, there are many metrics that have been developed in academia to measure model / AI fairness. These include 'equalised odds' and 'demographic parity'. Again as outlined above, the correct choice of metric will necessarily depend on the use case and may contain a subjective, ethical component that cannot readily be regulated for.

Other possible metrics to inform risk assessment might – depending on the use case – relate to:

- The degree of autonomy / automation vs human control over a system.
- Whether customers / individuals have some kind of ‘opt out’ available.
- The ease with which a system can be turned off.
- The impact of the use case on customers (if any), for example where the AI system could affect a customer’s product exposures.
- The level of reliance by the firm on a given system.

Under an outcomes-based approach, many metrics may not need to vary according to whether AI is involved. Complaint metrics, metrics regarding the accuracy and precision of credit risk assessments and lending decisions, metrics relating to timesaving, revenue generation or profit delivery, etc, would in large part be applicable to AI systems, simpler systems or human decision-making.

Although there are a range of metrics that can be useful *within* firms, it is challenging to design effective metrics to provide an overarching view of outcomes *across* the financial sector. For example, high numbers of objections or complaints could in fact be prompted by an organisation’s effective approach to fairness and efforts to be highly transparent with customers. While helpful at an organisational level, such statistics may be misleading as a tool to compare the performance of multiple organisations.

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

The discussion paper provides a comprehensive overview of the core requirements. As noted above, the suite of rules governing AI is already extensive, so care is needed to avoid creating duplicative new requirements.

Although not relating to the regulation of AI per se, we note that a key use case for AI in financial services is economic crime, such as transaction or client behaviour monitoring. This could be an area that would benefit from focused guidance in due course.

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?

At the present time we do not see significant regulatory barriers. Looking forward, however, we do see a risk that efforts to account for AI in regulation could lead to prescriptive rules being introduced over time that fail to accommodate differences between different use cases or models and techniques. Prescriptive rules could also create complexity and hamper innovation in the AI space. As outlined at the beginning of our response, a focus on regulatory principles and outcomes reduces these risks.

At the same time, an *absence* of regulation in a given domain or sector might push financial services firms to take a highly cautious approach. This could reduce innovation and partnerships with some vendors or across sectors, if firms lack confidence in dealing with them.

As set out under question 6, we do note that there is a degree of uncertainty in the sector over the best choice of fairness metrics and the interactions with the Equality Act, alongside

challenges for model testing and validation caused by GDPR requirements. But these are not areas for which the BoE and FCA have primary responsibility.

Given the complexity of layers of regulatory requirements across financial services, data protection, etc, understanding which existing rules can apply to AI use cases can be a barrier against AI uptake, particularly for SMEs.¹ A 'signposting' tool helping firms identify the different sectoral and horizontal rules that can apply to their AI use case could be beneficial. This would require collaboration with other authorities.

Lastly, we note that data localisation requirements impact on the effective use of AI and models by impacting the easing with which groups can share and combine their data. This is relevant, for example, in relation to Environment, Social and Governance efforts, where broad cross-sectoral datasets can be needed. We recognise of course that reducing barriers to data flows would require international effort and could not be solved by FCA and BoE alone.

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

Our members have already embedded specific applications of AI and ML into their processes. Their use is evolving and likely to continue to do so, bringing exciting opportunities to help serve customers better but with a concomitant risk of customer harm which firms and regulators must seek to avoid.

We recognise that there is already a considerable body of legal requirements in place that apply to the use of AI, in financial services specifically. These sit alongside cross-sectoral requirements which the DP helpfully and comprehensively lists. They show that regulators have so far wisely avoided taking premature steps that run the risk of altering the trajectory of innovation unduly.

Our view is that existing requirements, such as BCBS 239, GDPR and the PRA's proposed principles for model risk management, already provide an adequate framework for AI oversight and governance in financial services.

At the moment, this existing material sets expectations about the governance of AI. Firms' application of these governance expectations may well build with greater industry and supervisory experience. As it does so, AI users will continue to adapt their processes and oversight mechanisms to ensure that supervisory objectives continue to be met.

Overall, we do not believe that extra clarification of FS rules is needed at this time but recognise that as the experience and utilisation of AI deepens there may be a need for greater detail about the intersection between AI and regulatory requirements to close any gaps. That said, and building on our comments at the start of this response, we note that such clarifications would not necessarily take the form of 'AI guidance'. Rather, guidance should be targeted at the specific risks arising, for example focusing on certain types of models or system, such as opaque models, models that change autonomously while they are live, or systems that do not have a human directly in the loop.

In terms of identifying further areas where guidance will be valuable, we think it is important to first understand how the current regulatory framework fits together, in terms of which authority will be in charge of which issues and where there may be overlaps in practice. Although we

¹ This was identified by DCMS in its 2022 paper *Establishing a pro-innovation approach to regulating AI*. <https://www.gov.uk/government/publications/establishing-a-pro-innovation-approach-to-regulating-ai>

have a picture of the direction of travel from the Office for AI's policy paper published in the summer, this remains uncertain for now. We understand that the forthcoming AI whitepaper from the Office for AI should help provide clarity. (We nonetheless note under questions 6, 8 and 15 that *in due course* there could be benefit in guidance on AI discrimination, use of AI in economic crime detection or on AI risk factors).

We do also note that the GDPR rules for 'automated decision-making' are unclear at present. However, these provisions will be amended by the Data Protection and Digital Information Bill, to be supplemented with statutory instruments and ICO guidance, which should provide greater clarity in due course.

Q11: How could current regulation be simplified, strengthened and/or extended to better encompass AI and address potential risks and harms?

As we note above, any changes to the current suite of regulations and guidance should be made gradually as experience and understanding of the risks and opportunities that the application of AI brings increases. As noted above, any changes should not be duplicative or conflict with existing requirements. We favour an evolutionary approach based on experience and careful consideration of the challenges that deployment of AI brings, including concerns about safety, fairness, diversity, and privacy. Adding focused pieces of guidance over time that target specific new areas of uncertainty or risk will avoid duplication and unnecessary burdens.

An important feature of our members' development of new products and applications is the ability to 'experiment' in a sandbox or safe space before deploying a new product or system in the open market. Regulators' approach to AI should recognise this need for experimentation and permit firms to set their own principles in the context of their risk appetite before they release a new product to customers.

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

Depending on one's definition, AI and ML techniques are not necessarily wholly 'new' technologies per se, but a newer application of existing techniques, combined with more diverse data sets. As such, we believe that current governance structures have the elements necessary and are fit for purpose, on the whole, though with potential 'evolutions' to adapt to these new techniques and to connect up different governance components when necessary. Emerging AI governance approaches we see amongst some members currently include:

- new product management committees, concerned with fair customer outcomes and regulatory risk
- model risk management committees, concerned with data lineage and relevance and independent review and oversight
- cross-firm AI committees or data ethics committees that bring together a mix of interested and involved colleagues to address concerns in the context of live business decisions
- ensuring that the right skillsets and people are employed.

AI and ML techniques will undoubtedly evolve. The adaptation of regulation to this evolution should be a task shared jointly between industry and its regulators. Financial service regulators should not duplicate existing initiatives that are already being undertaken by other non-financial services specific regulatory bodies but rather import the experience of the wider AI ecosystem into their own expectations where relevant. This will require a degree of horizon scanning, which economy-wide organisations such as the CDEI and the Alan Turing Institute

would be well placed to coordinate. Similarly financial services firms should be alive to how other sectors and jurisdictions are addressing the challenge; monitoring external ideas will help ensure that best practices are kept pace with.

Our members often use SaaS providers to combine multiple AI components from third-parties or themselves utilise and combine AI-enabled solutions. We should not overlook the importance of 3rd party providers contributing to the debate but note that this is already an intense area of work by authorities.

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?

AI is not a technological process in its own right, but a means to an end that has the potential to improve customer experience and mitigate – for example – conduct, prudential and competition risk. A range of different individuals will be accountable for their firm’s management of these risks. We do not think it is realistic to have a single individual responsible for all ‘AI risks’, given how diverse these are.

Responsibility will, in our view, be more effective and clearer when focused on specific risk types, rather than specific technologies. We note that the SMF responsible for model risk under the draft CP 6/22 proposals might often take a leading role but also that, depending on how the firm’s internal governance and business are structured, different functions might be appropriately involved; for example, a number of individuals at a level below SMF may have direct AI responsibilities.

A core feature of the Senior Managers and Certification Regime (SM&CR) is to ensure that business-line and function-aligned SMFs are held accountable for actions within their remit. Allocating Prescribed Responsibility for AI to a single SMF would undermine this principle, given that AI will be pervasive across organisations, developed and implemented by different business teams and functions under the responsibility of a number of different SMFs.

Indeed, we have generally advised against creating new Prescribed Responsibilities so we do not think that a new Prescribed Responsibility for AI would provide greater assurance.

Q14: Would further guidance on how to interpret the ‘reasonable steps’ element of the SM&CR in an AI context be helpful?

There is currently little regulatory material about what constitutes ‘reasonable steps’. It is our understanding that, at a high level, ‘reasonable steps’ encompasses behaving with integrity, appropriate delegation, understanding of the senior managers’ business area, and complying with law, rules, ethics and legal obligations. The detail will differ depending on the size and type of the organisation, but this is likely to engage other regulatory duties, such as the Consumer Duty where applicable.

We do not believe that ‘reasonable steps’ for AI & ML requires additional guidance to these overarching common-sense tests, nor should they be applied in a different manner for AI relative to any other matter.

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?

We agree that a more coordinated approach between supervisory authorities, both within and beyond the financial services sector, on taxonomies and terminology related to AI would be helpful and would likely lead to increased legal certainty for firms.

We also believe there may be particular merit in greater alignment in terminology on AI in relation to customer-facing communications to facilitate consistency and consumer understanding. To deploy AI, firms will need to be able to explain transparently and accessibly how an algorithm or model defines similarities between customers, why certain differences between two prospects may justify different treatments, and why similar customers may get different explanations about the AI. Importantly, explanations of AI decisions may improve fairness and increase stakeholders' trust.

In relation to data governance, we consider that a single unified approach is unlikely to be feasible, as the data inputs will vary greatly depending on the use case. There might, however, be benefit in developing risk indicators / a tiered risk approach.

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?

Authorities can ensure that, as the regime evolves, we retain clarity as to where responsibilities lie. The EU's AI Act has tried to split responsibility between providers and users of AI but, in practice, this is likely to be difficult to resolve. For example, where a 'user' buys in a vendor model but then enhances it through further training on the user firm's data, the roles will become blurred.

Authorities might also explore options for techniques for identifying and addressing unfair bias or discrimination, such as the use of synthetic data.

Lastly, we encourage authorities to coordinate on AI regulation issues. As noted above in our response, as well as in the discussion document, there are overlaps between rules which will need to be managed in order to ensure a level playing field, avoid gaps and minimise regulatory arbitrage, particularly in relation to use cases that are cross-sectoral, or AI products sold into multiple sectors. There is a risk that a lack of coordination could lead to an increasingly confusing overall regulatory picture, which could inhibit beneficial AI uptake in the sector.

Similarly, although we support the more flexible (though still provisional and emerging) UK approach over the EU's more rigid statutory framework, the reality is that many firms and their suppliers will need to interact with both jurisdictions' regimes. As such, ways to collaborate and streamline should be explored with EU stakeholders, within the bounds of what is feasible under the different approaches.

The AI assurance ecosystem also has a role to play across sectors, for example by helping firms to leverage third party solutions. This would be another area that could benefit from coordination among sectoral authorities, for example in support of the ongoing work of the CDEI.

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?

We generally consider that the focus of regulatory authorities should be on setting the high-level outcomes and principles for firms, supplemented by guidance when appropriate / necessary. The draft framework under CP6/22 or under the Federal Reserve's SR 11-7 are examples of this approach in the context of model risk management, setting boundaries but also allowing leeway to firms on implementation and management.

We are supportive of the role of voluntary standards as a tool available to industry. Voluntary global standards, such as those being developed by ISO/IEC SC 42, and guideline frameworks, such as NIST the AI Risk Model Management Framework or Veritas Consortium assessment methodologies, can help businesses to use AI responsibly, which is welcome.

However, we recommend caution in relying on, or referring to, AI standards in guidance and regulation. This is because:

- Most AI standards are still under development and the industry requires time to evaluate these upon publication.
- Many standards will – by their nature – be applicable to some businesses and not to others. Businesses should also have the flexibility to address AI challenges in other ways, if they believe this will lead to the best outcome for their business, consumers and the ecosystem or if their AI challenge is not covered by standards.
- Standards are reviewed every three to five years, and new state of the art solutions could be available for businesses in the meantime to address AI challenges more effectively.

As such, referencing standards in regulation could stifle innovation and might require regulation to be reviewed more frequently.

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?

Building on our support for principles-based and outcomes-focused regulation, we highlight that this kind of approach facilitates international interoperability. This facilitates economies of scale by firms with international operations or international customers, clients or suppliers, reducing costs for firms and ultimately consumers, and drives greater consistency.

Where different jurisdictions have highly prescriptive rules, this leads invariably to differences that require siloed policies and procedures, creating duplication and complexity, and greater difficulty for firms 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?

Please see our comments on prescriptive regulatory approaches throughout this response.

ENDS