Holistic AI Insight Paper

Regulating AI in Financial Services

Holistic AI responses to the regulators' proposals

October 2022

Holistic AI

HOLISTICAI.COM

CONTENTS

3 Executive Summary

5 Background

7 Overview of the Discussion Paper

11 Analysis and Recommendations

19 Conclusion

KEY TAKEAWAYS

- The UK Government's approach to AI regulation is to empower regulators to develop context-specific rules and guidance for their respective sectors.
- In October 2022, the Bank of England (BoE), the Prudential Regulation Authority (PRA) and the Financial Conduct Authority (FCA) published a Discussion Paper on regulating AI and machine learning (ML) in UK financial services.
- The above mentioned 'Discussion Paper' emphasises that AI poses novel risks and challenges, whilst also amplifying existing risks. A proactive approach to AI Risk Management is therefore required.
- The Discussion Paper clarifies the ways in which the existing regulatory framework applies to AI use (e.g., data protection, consumer protection, equalities law etc.).
- Under this framework, financial services firms have extensive risk management obligations. These processes and systems must be reviewed, adapted and revamped to the address the novel risks and challenges of AI use.
- Furthermore, additional guidance and regulation will be required, to protect consumers and support firms in navigating issues like AI procurement, independent algorithmic auditing, mitigating bias and discrimination, and promoting transparent and explainable AI.
- AI regulation is a top priority issue for the UK Government and regulators. Firms should follow these developments closely, and take proactive and practical steps to understand and manage their AI risks.

1 EXECUTIVE SUMMARY

This Holistic AI Insight Paper provides an overview and analysis of the Discussion Paper on regulating AI and machine learning (ML) in financial services, which was published in October 2022 by the Bank of England (BoE), the Prudential Regulation Authority (PRA) and the Financial Conduct Authority (FCA).

It is informed by Holistic AI's industry leading AI Risk Management work. Holistic AI has extensive practical experience in this area, having reviewed and audited over 100 enterprise AI projects, covering more than 20,000 different algorithms. Our clients and partners include Fortune 500 corporations, SMEs, governments and regulators.

There is widespread agreement among the UK Government, regulators and parliament that AI poses novel risks, which have the potential to cause significant harm.

72% of financial services firms use or develop AI / ML, with adoption rapidly increasing across all business areas. The insurance sector is experiencing the largest increase in AI / ML use.

The financial regulators' Discussion Paper comes shortly after the UK Government announced that regulators will be empowered to develop pro-innovation, proportionate and context-specific rules and guidance for their repsective sectors.

The Discussion Paper outlines the main benefits and risks of AI in financial services. The benefits include increased efficiences, more accurate decision-making, enhanced fraud detection, and more personalised products and services. Risks and harms include bias and discrimination, inaccurate predictions, a lack of explainability, and risks to financial and market stability.

The Discussion Paper also clarifies how the existing regulatory framework applies to AI, seeking feedback on whether it should be updated. It focuses on areas including governance and risk management, equalities law, data protection, senior management accountability (SM&CR), and outsourcing and third-party risk management.

The regulators are sending a clear message: firms must review, adapt and revamp their risk management processes and systems in the context of AI. As many AI risks are novel, the approach to managing them must be adapted to reflect that. This work should be prioritised, given the increasingly widespread and material use of AI.

Financial services firms should integrate the core elements of AI Risk Management into their approach. This includes ongoing and dynamic monitoring of AI systems in real-time, given that AI systems continuously learn and evolve. For firms deploying AI across the business, this ongoing monitoring is virtually impossible without a scalable and automated AI inventory manageemnt solution.

Furthermore, AI risks can emerge from issues at each stage of the AI lifecycle. The earlier in the lifecycle risks are surfaced, the better. This means that technical assessments and tests need to be performed, and risk mitigations deployed, throughout the AI lifecycle.

There are several areas where additional guidance and new regulations will be required, to protect consumers and markets, and to support firms in managing AI risks.

For example, the regulators should clarify their position on the independent auditing and review of AI systems, which they ostensibly endorse as an important part of the solution.

Also, firms would benefit from additional guidance regarding how they should monitor and maintain high levels of AI performance, how they should approach procuring AI systems from third party vendors, and what constitutes 'reasonable steps' for risk management at each stage of the AI lifecycle.

Other jurisdictions (e.g., U.S. states, EU, Canada) have proposed or enacted regulations requiring the auditing of AI systems for bias and the use of of representative datasets to train AI. It is worth considering whether a similar approach should be adopted in the UK.

Finally, AI risks cannot be adequately managed without robust governance and accountability. There needs to be a senior management function and/or individuals who are responsible for AI Risk Management.

AI regulation is a top priority issue for both the government and financial services regulators. Firms should follow these developments closely, and take proactive and practical steps to understand and manage their AI risks.

2 BACKGROUND

2.1 Discussion Paper on regulation of AI in financial services

On 11 October 2022, the Bank of England (BoE), the Prudential Regulation Authority (PRA) and the Financial Conduct Authority (FCA) jointly published a Discussion Paper on regulating artificial intelligence (AI) and machine learning (ML) in UK financial services.

In recent years, regulators, government departments and parliamentary committees have published extensively on AI. There is widespread agreement that AI regulation is a top priority issue, because AI poses novel risks, which have the potential to cause significant harm.

The purpose of the Discussion Paper is to provide insight on the main benefits and risks of AI in financial services, and explore whether the existing regulatory framework should be updated.

2.2 The UK Government approach to AI regulation

The UK Government published its AI Regulation Policy Paper in July 2022, which outlined plans to establish a pro-innovation, proportionate and context-specific framework.

The Government's preference is to pursue targeted, non-statutory options (e.g., soft law and guidance).

The UK's priority is to avoid stifling innovation and growth with overly prescriptive and burdensome regulations, whilst simultaneously "keeping people safe and secure" via responsible AI use.

The government is empowering regulators to take the lead, by developing sector-specific rules and guidance, which implement the following 'cross-sectoral' principles: safety, security, transparency, fairness, legal responsibility and governance, and clarity of redress and contestability.

A government White Paper on AI regulation is planned for later this year.

2.3 UK financial authorities and AI regulation

The precise nature of AI regulation in the financial services sector will thus be jointly determined by the BoE (responsible for monetary and financial stability and market integrity), the PRA (supervises major

financial institutios) and the FCA (upholds consumers protection, effective competition and market functioning).

This Discussion Paper is not their first intervention on AI. In June 2022, the PRA published a Consultation Paper, which proposed a set of principles for banks' Model Risk Management, including: model identification and risk classification; governance; model development, implementation and use; independent model validation; and model risk mitigants.

Also, the FCA is part of the Digital Regulation Cooperation Forum (DCRF), along with the Competition and Markets Authority (CMA), the Information Commissioner's Office (ICO) and the Office for Communications (Ofcom).

The DCRF's priorities for 2022-23 include 'supporting improvements in algorithmic transparency', 'improving capabilities in algorithmic auditing' and 'promoting transparency in algorithmic procurement'.

How AI is used in financial services

Algorithmic trading

Insurance underwriting and claims management

Automated mortgage and loan decisions

Regulatory capital modelling

Fraud and AML detection

Personalisation of products and services

3 OVERVIEW OF THE DISCUSSION PAPER

The Discussion Paper outlines the main benefits, risks and harms of AI in financial services. It also clarifies how the the existing regulatory framework applies to AI, and strongly indicates that new guidance and regulation will be required.

3.1 AI and machine learning in UK financial services

The BoE and FCA's recent survey on AI / ML in large UK financial services firms found that:

- 72% of firms use or develop ML
- ML adoption is rapidly increasing across all business areas
- The insurance sector is experiencing the largest increase in ML use
- Most firms do not see current ML use as high risk

3.2 Benefits, risks and harms of AI in financial services

The Discussion Paper outlines the main benefits, risks and harms of AI in UK financial services.

The benefits include reduced costs and processing times (e.g., insurance claims management), more accurate decision-making (e.g., credit default risk prediction), enhanced fraud detection (e.g., AML screening) and greater personalisation (e.g., investment products).

The risks and harms include bias and discrimination (e.g., insurance pricing), flawed predictions (e.g., inaccurate capital modelling), a lack of explainability and interpretability (e.g., inability to understand transactions flagged as fraudelent) and risks to financial and market stability (e.g., algorithmic 'herding' causing flash bubbles and crashes).

The Discussion Paper emphasises the importance of taking an AI lifecycle approach, meaning that different actions should be taken at each stage of the AI lifecycle (pre-deployment, deployment, and recovery and redress).

It argues arguing that the primary drivers of AI risk stem from:

- **Data** (training, testing and validation data)
  - Risk factors: unrepresentative data; historical biases; lack of quality control; poor ingestion; significant outliers or noise; lack of monitoring and reporting.
- **Model**
  - Risk factors: data-related risks; inappropriate model choice; model design and construction errors; lack of explainability; model or concept drift; degrading performance; unexpected behaviour.
- **Governance**
  - Risk factors: unchecked autonomous decision-making; lack of human in the loop; lack of accountability; absence of well defined roles and responsibilities; skill and expertise gaps.

The regulators argue that issues at each stage of the AI lifecycle can create risks, resulting in harm to consumers, firms and financial markets.

3.3 Clarification of how the existing regulatory framework applies to AI

The Discussion Paper clarifies and explains the myriad ways in which the UK's existing legal and regulatory framework for financial services applies to AI.

Given the highly regulated nature of the sector, this is an important resource for firms seeking to understand their obligations in the context of AI development and use.

The Discussion Paper's intention is to generate feedback on whether the existing regulatory framework is sufficient for regulating AI, or whether new regulations are needed.

The table below highlights the Discussion Paper's assessment of the primary ways in which the exsting regulatory framework applies to AI:

Area of concern (lead regulator)	Relevant law, regulation or guidance	Relevance for AI
Consumer protection (FCA)	**Equality Act 2010** FCA's **Principles for Business** and the **'Consumer Duty'**	Algorithmic bias discrimination (e.g., in pricing and customer risk segmentation).
Data (ICO, PRA and FCA)	**UK GDPR and Data Protection Act 2018** **Principles for effective risk data aggregation and risk reporting** **The Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017**	Rules governing the collection, processing and use of data used to to train, test and validate AI systems (e.g., data minimisation; data security; accuracy requirements). Restrictions on automated decision-making and profiling.
Governance and risk management (PRA and FCA)	**FCA Principles for Business** **PRA Rulebook** **Corporate governance: Board responsibilities** (PRA SS5/16)	Firms must establish and implement adequate risk management and governance systems and processes (e.g., defined roles and responsibilities; sufficient technical skills and expertise at board and senior levels). These systems and processes must be adapted to address the novel risks of AI.

Model risk management (PRA)	**Proposed model risk management (MRM) principles for banks**	Once enacted, firms will have to apply these MRM to the governance of AI and ML models which inform key business decisions.
Accountability of senior management (PRA and FCA)	**Senior Managers and Certification Regime** (SM&CR)	The SM&CR aims to protect consumers and strengthen market integrity by enhancing senior leadership accountability. Although there is currently no assigned Senior Management Function (SMF) for AI, its use in different business areas will fall within the scope of different SMFs.
Outsourcing and third-party risk management (PRA and FCA)	**FCA Handbook** **Outsourcing and third party risk management** (PRA SS2/21)	Rules relating to outsourcing and third party risk management are relevant for AI procurement.

4 ANALYSIS

This section provides analysis and recommendations on the Discussion Paper's key contributions and proposals.

4.1 AI Risk Management: a novel approach

Discussion Paper, 4.45: Principle 3 of the FCA Principles for Businesses requires that '[a] firm must take reasonable care to organise and control its affairs responsibly and effectively, with adequate risk management systems'.

By outlining firms' extensive risk management obligations -- and simultaneously emphasising that AI both amplifies existing risks and poses new risks -- the Discussion Paper relays a clear message: firms should review, adapt and revamp their approach to risk management in the context of AI development and use. Given the increasingly widespread and material use of AI across different business areas, firms must prioritise this work.

Firms should be aware that although risk management principles and obligations apply to AI, regulators are unlikely to deem 'business as usual' approaches as adequate, given the novel risks and challenges which they have highlighted. Put simply, as many AI risks are new, the approach to managing them must be adapted to reflect that.

To ensure compliance with existing risk management obligations in the context of AI use, firms should integrate the core elements of AI Risk Management into their approach, including:

- Ongoing and dynamic monitoring of AI systems and their outputs in real-time
- Scalable and automated AI inventory management
- Technical and engineering-based assessment, testing and risk mitigation
- Updated policies and processes, including governance and accountability frameworks (with well defined roles and responsibilities)

4.2 AI lifecycle and continus monitoring

Discussion Paper, 4.59: One useful approach to understanding firms' obligations is to look at them from the perspective of the AI lifecycle. For example:

- *pre-deployment: how should the quality of training data be assessed? How should AI models be tested before live deployment? Should AI models be 'compliant by design'? Who is the accountable SMF? Identification of and allocation of responsibility for new risks presented by AI.*
- *deployment: how should the performance of live AI systems be monitored? What safeguards should be introduced to monitor for, detect and stop potential harm e.g., kill-switch mechanisms?*
- *recovery and redress: if an AI system's performance leads to crystallised risks, should firms be required or expected to 'undo the damage' by (i) reversing decisions made by the model (where possible and appropriate); and/or (ii) compensating any relevant external parties who suffered damage as a result?*

The Discussion Paper's emphasis on the AI lifecylce is appropriate and well-founded.

The underlying premise that AI poses novel risks, necessitating a distinct approach to AI Risk Management, is correct. The importance of ongoing, real-time monitoring of AI systems is a key factor which differentiates AI Risk Management from more traditional approaches to risk management.

Given that AI systems continuously learn, adapt and evolve, they need to be carefully monitored on an ongoing basis. The higher-risk the system, the more frequently it should be assessed and reviewed. The fact that AI performance tends to decay over time reinforces the importance of this. For organisations deploying many AI systems, this monitoring is virtually impossible without some form of scalable and automated solution.

AI risks can emerge at or stem from any stage of the AI lifecycle, meaning every stage is relevant. The earlier in the lifecycle risks are surfaced and mitigated, the better.

It would be useful for additional guidance to be produced on the concrete steps which firms should take at each stage of the AI lifecycle. Also, future guidance could break down the different stages of the lifecylce in a more granular and accessible way.

4.3 AI inventory management

Discussion Paper, 4.68: Operational resilience: firms and FMIs should set an impact tolerance for disruption for each of those important business services that involve AI, and ensure they are able to remain within their impact tolerances for each important business service in the event of a severe (or in the case of FMIs, extreme), but plausible disruption.

Many of the actions that firms will be expected to take, like maintaining operational resilience, are predicated on the underlying expectation that firms should be cognizant of all the AI systems they use across the business. However, most organisations lack this deep understanding of their 'AI inventory'.

To be compliant, firms require deep understanding of their AI inventory. This generates actionable insights which inform AI Risk Management strategies. Once a scalable and automated system for monitoring the AI inventory is established, AI outputs can be monitored, and interventions can be employed to mitigate risks.

For example, to meet obligations to maintain operational resilience in the context of AI use, firms must know:

- How many AI systems are being used?
- Which business areas are they being used in?
- How dependent are those business areas on AI?
- What is the overall risk profile (and corresponding likelihood of disruption) of the AI use in specific business areas?

Firms can only achieve this understanding with a scalable and automated AI inventory management solution.

The PRA's position in its model risk management Consultation Paper is that firms should 'maintain a model inventory' and 'classify models into risk tiers'.

Given the cross-cutting importance of this issue, regulators should consider whether firms should be mandated to maintain an inventory of all AI and ML systems they use. At a minimum, guidance or clarificaiton on best practice for AI inventory management should be issued.

4.4 Independent auditing and review of AI systems

Discussion Paper, 4.42: The validation and independent review of an AI model is important in order to ensure an objective view is given on the model, inclusive of the way in which it is developed and that it is suitable for the intended purpose.

Independent auditing and review of AI systems is deemed an important component of managing AI risks. Firms should explore options in this domain when designing and establishing their AI Risk Management frameworks.
This is a priority issue for both the FCA (as highlighted in the **DRCF's 2022-23 workplan**) and the UK Government (see the Centre for Data Ethics and Innovation (CDEI) workstream on **AI Assurance**).

Furthermore, the Discussion Paper otensibly endorses the International Organisation of Securities Commissions's (IOSCO) stance that 'regulators should require firms to adequately test and monitor their algorithms to validate the results of an AI technique on a continuous basis'. Many firms would likely struggle to do the without external support (i.e., bespoke technical expertise and independent auditing).

It would be useful for the regulators to clarify their position on the independent auditing, review and testing of AI systems. It would also be useful for the regulators to support future work to build empirical evidence and case studies in this area, to inform best practice.

4.5 Insufficient focus on robustness and efficacy risks

Discussion Paper, 3.6: Poor AI model performance may result from data-related risks but also from a range of model-related risks. These could include inappropriate model choices, errors in the model design or construction, lack of explainability, unexpected behaviour, unintended consequences, degradation in model performance, model or concept drift, and more.

Current ML use cases **are not viewed** as high risk by most financial services firms. Also, the Discussion Paper does not extensively focus on how firms should address the risks stemming from poor or unexpected model performance. This suggests that insufficient attention is being paid to the risks of AI performance issues, at both the firm and regulator level.

Model performance (i.e., whether and how well an AI system works) is usually the most significant risk factor in terms of real-world impact. Despite this, organisations often lack a standardised and comprehensive approach for assessing, monitoring and improving AI performance.

The performance of an AI system can be measured in terms of both robustness and efficacy. Robustness refers to how stable the system's performance is when changes or attacks occur, and efficacy refers to the system's performance relative to its intended purpose. Some contexts require higher levels of efficacy than others (e.g., regulatory capital modelling).

The testing of an AI system against robustness and efficacy metrics reveals the system's overall performance level, and consequently, whether its outputs can be trusted as accurate and safe in a range of different and unpredictable scenarios.

Common AI performance issues are:

- The required performance level (%) set by the system's developer is too low for the context in which the system is deployed.
- The system's performance fluctuates significantly when presented with unseen data.
- The system produces different outputs even when given the same set of inputs, meaning it is not reproducible.

Such performance related issues undermine the accuracy, reliability and safety of AI systems. Fortunately, once identified, there are a range of technical mitigations which can be employed to improve a system's performance, like 'adversarial training' (training systems to withstand adversarial attacks) and the retraining of a model on new data sets to halt performance decay.

Although the Discussion Paper alludes to the risks of model performance issues and inaccurate AI decision-making, firms would benefit from additional guidance and clarity regarding how they should monitor and maintain high levels of model performance on an ongoing basis. Awareness of this issue, and the state-of-the-art in the field, is relaitvely low.

4.6 'Reasonable steps'

Discussion Paper, 4.60: The concept of 'reasonable steps' is a core element of the SM&CR. One of the areas that could benefit the most from further discussion is what may constitute reasonable steps in an AI context and how, if at all, these steps differ from the reasonable steps that SMFs are generally required to take. A particularly useful approach could be to consider what may constitute reasonable steps at each successive stage of the lifecycle of a typical AI system.

The concept of 'reasonable steps' is pertinent for AI Risk Management, as there is an established and growing body of knowledge on how to detect, verify, mitigate and prevent AI risks.

It should be considered 'reasonable' for firms, many of which are critical to the stability of the financial system, to keep up with industry and engineering best practice and incorproate this into their risk management work.

The GDPR, for example, obliges organisations to implement state-of-the-art cybersecurity measures, which refect industry consenus on best practice at the time of implementation.

There are a range of state-of-the-art tests which can be performed on the datasets and algorithmic models which AI systems are comprised of. These tests can identify issues relating to robustness, performance, bias, transparency, privacy, and

other responsible AI verticals. Furthermore, there are technical mitigations which can be employed to address these issues. The earlier in the AI lifecylce these tests and assessments are performed, the more likely it is that any issues are identified, mitigated and resolved.

It would thus be useful for the financial regulators to produce additional guidance on what constitutes 'reasonable steps' at each stage of the AI lifecycle, with reference to the wide range of technical and engineering-based tests and mitigations that exist. It would also be useful to clarify whether firms are obliged to implement industry consenus on best practice with respect to AI Risk Management.

4.7 AI procurement

Discussion Paper, 3.23: a further key challenge for firms lies in their ability to monitor operations and risk management activities that take place outside their organisations at third parties. Increased reliance on third parties, often outside the regulatory perimeter, for datasets, AI algorithms, and other IT outsourcing (such as cloud computing) may amplify systemic risks.
The regulators are right to highlight that the procurement of AI from third parties can amplify risk.

Most enterprises purchase the AI tools they use from technology vendors. Asymetries of information and technical expertise means that the customer often lacks the capabilites to ensure that the AI systems they are purchasing from the supplier are fit for purpose and will perform as intended.

Financial services firms would benefit from additional guidance on AI procurement, given the centrality of the issue. This could build upon previous government **Guidelines for AI Procurement**, which is relatively high-level.

Key questions which such guidance could cover are:

- what due diligence firms should conduct prior to procuring AI
- which information and instructions should suppliers provide
- which tests and assessments should be undertaken before deployment of a procured AI system tool

4.8 Bias and discrimination

Discussion Paper, 4.21: The Consumer Duty also addresses discrimination harms by requiring firms to consider the diverse needs of their customers, including the fair treatment of customers with characteristics of vulnerability and those with protected characteristics. Firms will be required to monitor the outcomes their customers receive in practice and take action if they identify particular groups of customers are getting poor outcomes.

The Discussion Paper focuses extensively on the risks of bias and discrimination, including how firms are obliged to ensure that their systems and processes are

compliant with equalities legislation and the Consumer Duty. However, there is little practical information or guidance on exactly how firms should monitor, detect and prevent bias and discriminatory algorithmic decision-making.

Other jurisdictions have imposed specific requirements aimed at highlighting, detecting and preventing bias. For example, **New York City Local Law 144** mandates independent bias audits of 'automated employment decision tools' which are used by employers or employment agencies to employ candidates or promote employees residing in New York City. The results of the bias audits must be published. Also, **Colorado Senate Bill 21-169** prohibits 'unfair discrimination' in insurance practices and mandates testing, corrective action and risk management of AI and algorithmic systems.

The best way to improve outcomes for consumers is to minimise the risk of biased decision-making at the earlier stages of the AI lifecycle (i.e., pre-deployment). However, there are currently no explicit requirements to test or audit systems for bias.

Discussion Paper, 2.15: AI may pick up bias within datasets and may not perform as intended when exposed to issues excluded from the training/testing data, so new data quality metrics like representativeness and completeness may be needed.

The proposed **EU AI Act** requires training, testing and validation data sets to be representative and relevant for the intended use case.

Given that unrepresentative training data is one of the main ways in which AI amplifies and perpetuates historical biases, it is worth considering whether a similar approach should be adopted in the UK.

4.9 Explainability and interpretability

Discussion Paper, 4.16: Certain AI-derived price-discrimination strategies could breach the requirements if they result in poor outcomes for groups of retail customers. As such, firms should be able to monitor, explain, and justify if their AI models result in differences in price and value for different cohorts of customers.

Explainability and interpretability is crucial, as it enables different stakeholders to understand how and why an AI system makes decisions. This information can be used to intervene and take corrective action if there are any issues, understand the risk profile of the AI system, and provide meaningful and easily understandable explanations to external stakeholders (e.g., the board, shareholders, consumers, regulators) regarding what AI is being used, how it works, and the significance and impact of the outputs it generates.

It would be useful for regulators to clarify precisely how firms should communicate with their customers or other stakeholders, with respect to justifying and explaining how decisions were made and which factors were

taken into account. Such guidance could determine benchmarks for easily understandable and accessible explanations for different stakeholders in different contexts.

4.10 Accountability and governance

Discussion Paper, 4.51: PRA-authorised SM&CR banking and insurance firms and FCA-authorised enhanced scope SM&CR firms must ensure that one or more of their SMF managers have overall responsibility for each of the activities, business areas, and management functions of the firm. That means any use of AI in relation to an activity, business area, or management function of a firm would fall within the scope of a SMF manager's responsibilities. [...] 4.56: Looking ahead, there is a question as to whether there should be a dedicated SMF and/or a Prescribed Responsibility (PR) for AI under the SM&CR.

The Discussion Paper cites IOSCO guidance which recommends that 'regulators should consider requiring firms to have designated senior management responsible for the oversight of AI development, testing, deployment, monitoring, and controls.' This a key area of inquiry for the regulators, who are actively considering amending the SM&CR framework to assign a dedicated SMF for AI. Firms should prepare for this eventuality.

AI risks cannot be adequately managed without robust governance and accountability frameworks. There needs to be a senior management function and/or individuals who are responsible for AI Risk Management.

Given the complexity of this area, the rapid pace at which it is evolving, the increasing centrality of AI to core business functions, operations and decision-making, and the potential risks AI poses the integrity and stability of the financial markets, regulators should not wait until AI reaches higher levels of maturity (and correspondingly poses higher levels of risk), before requiring a dedicated SMF and/or PR for AI under the SM&CR.

5 CONCLUSION

Governments and regulators worldwide are focusing on how the risks and harms of AI use can be mitigated, whilst simultaneously promoting innovation and growth.

The consensus is that AI poses novel risks, requiring firms to be proactive in taking practical steps to identify, assess, mitigate and prevent them.

The EU AI Act will impose a sweeping set of mandatory requirements on the 'providers' and 'users' of high risk AI systems. Canada's AI and Data Act requires the designers and developers of 'high-impact' AI systems to establish risk management measures.

In the U.S., the White House recently published its Blueprint for an AI Bill of Rights, with proposals for a US Algorithmic Accountability Act having been tabled earlier this year. Different states and jurisdictions, such as New York City, Colorado and Illinois have enacted laws regulating the use of AI in HR, insurance and other sectors.

In the UK, financial services regulators are sending a clear message to the market: they are prioritising the issue of AI regulation and are actively considering introducing new guidance and rules.

Firms should follow these developments closely, and take proactive and practical steps to understand and manage their AI risks.

ADDITIONAL READING

Artificial Intelligence and Machine Learning – BoE, PRA and FCA Discussion Paper (2022)

Machine learning in UK financial services – BoE and FCA survey report (2022)

AI Public-Private Forum final report – BoE (2022)

The use of artificial intelligence and machine learning by market intermediaries and asset managers – IOSCO Final Report (2021)

AI in Financial Services – Alan Turing Institute (2021)

Model risk management principles for banks – PRA Consulation Paper (2022)

Appendices to CP6/22 – Model risk management principles for banks (2022)

The Impact of Big Data and Artificial Intelligence (AI) in the Insurance Sector – OECD (2020)

AI Governance Principles – EIOPA (2021)

Establishing a pro-innovation approach to regulating AI – UK Government (2022)

Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector – Monetary Authority of Singapore

GOT QUESTIONS OR WANT TO SCHEDULE A CHAT? CONTACT US AT

holisticai.com we@holisticai.com

Holistic AI