AI in financial services

Are you meeting the regulators' expectations?

December 2019

pwc

Contents

AI in financial services – Are you meeting the regulators' expectations?

> "Artificial Intelligence and machine learning... [have] the potential to yield enormous benefits for households and businesses by opening up new lines of credit, providing greater choice, better-targeted products and keener pricing."
>
> **Mark Carney, Governor of the BoE**

The financial services (FS) industry is going through a period of profound change and disruption. Technology is providing the means for firms to reimagine the way in which they operate and interact with their customers, suppliers and employees. One significant area of development is the utilisation of artificial intelligence (AI) and machine learning (ML). The majority (56%) of global financial services executives we surveyed see AI as set to transform the way in which their services are delivered within the next two years[1]. But as with any change, regulatory interest in, and scrutiny of, AI is growing – and firms are facing new and exacerbated risks such as explainability, bias, governance, accountability and data protection across the AI lifecycle. In this report we explore what firms need to do to meet the regulators' expectations on AI and ML – and why it's crucial that they do so.

Introduction

A burning issue for FS firms adopting AI and ML technologies, highlighted by the BoE and FCA in an October 2019 survey they conducted, is a need for greater regulatory clarity[2]. Given the broad adoption of technology we have already seen in the sector, it is no surprise that senior regulators at both the BoE and FCA have started to set out some of their key expectations around the use of these technologies, especially AI and ML. The international debate on the regulatory treatment of AI is also developing with the European Commission, IOSCO, Financial Stability Board, Monetary Authority of Singapore and the US Treasury Department all sharing their views on the issue.

Although a regulatory consensus on the treatment of AI and other technologies has not yet emerged, our research shows that firms are already on the front foot in rolling out extensive work programmes to utilise their power. Over 44% of UK FS organisations have already embedded FinTech fully into their strategic operating model (48% have globally), while 27% of FS organisations have also incorporated emerging tech into the products and services they sell (37% globally). Of course, being a first-mover has its advantages, but it also runs the risk of misalignment with regulatory expectations. Therefore, understanding regulatory expectations around the use of AI and factoring them into implementation and monitoring programmes is a vital first step if firms are to exploit the potential of AI and ML technology, in a sustainable and operationally resilient manner.

---

[1] PwC Global Fintech Report 2019

[2] Machine learning in UK financial services, BoE and FCA

While our report explores five crucial areas for firms to consider right now, it is important to note that regulators in the UK are still exploring what the future of regulation could look like. This could take the form of a principles and outcomes-based approach, leveraging existing regulation, or the construction of AI-specific legislation. It also remains to be seen whether the agenda will be dictated primarily at a national level, or if a greater emphasis is placed on international collaboration. For the time being though the FCA and BoE remain technology agnostic. This means firms will be obliged to treat customers fairly, establish robust governance arrangements and manage risk irrespective of the mechanism through which they provide services.

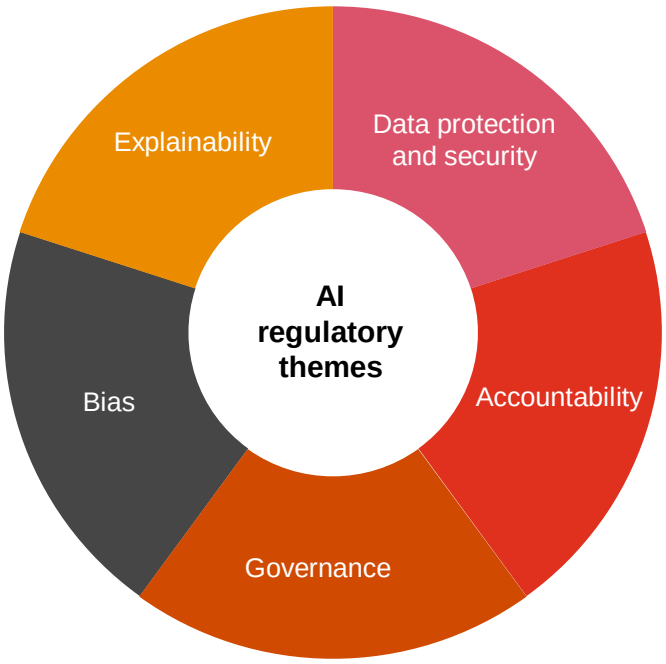

AI regulatory themes

1. Explainability

Explainability relates to how AI systems reach outcomes, and how the rationale for these outcomes can be explained. Better explainability of models should help to overcome issues around the 'black box' – that is, the processes that occur between the initial input and end output, which cannot always be easily understood and summarised. Ultimately, understanding how and why a decision is made is often as critical as the accuracy of the result – particularly when the decision has material consequences for consumers or broader financial stability.

84%

of CEOs agree that AI-based decisions need to be explainable in order to be trusted[3]

Regulators are assessing how firms can balance the trade-off between harnessing AI to make effective decisions while also being able to explain them to stakeholders. The FCA has signalled that firms should focus on achieving 'sufficient interpretability', essentially a compromise between ensuring AI functions as intended while also recognising the need for firms to make sense of the main drivers of the decision-making process. Firms will also need to consider their existing obligations under GDPR. Solely automated decisions which have a 'significant' or 'legal' effect on an individual are prohibited (unless explicitly authorised by consent). Where they

[3] 22nd PwC Global CEO Survey

are permitted, individuals should have recourse to have the outcome of any decision reached explained to them.

The ICO and The Alan Turing Institute are also collaborating to create practical guidance for firms to explain AI decisions to affected individuals through Project ExplAIn. The final guidance is due to be published in early 2020, but an interim report published on 3 June 2019 has already highlighted a number of key findings. These include the lack of a one-size-fits-all approach to explanations, the need for board-level buy-in on explaining AI decisions, and the importance of a standardised approach to internal accountability to help assign responsibility for explainable AI decision-making.

Key considerations for firms

- Firms will need to determine the level at which explainability is pitched, with reference to the type of decision informed by the algorithm and the potential impact on the stakeholder concerned. For example, an automated decision to refuse a mortgage application will need to be explained in more simple and straightforward terms than is likely to be possible in practice for an FX trading algorithm. Secondly, while a decision made by an AI system may be explainable to a firm's Chief Digital or Data Officer, would a retail customer understand the implications?
- Firms should also consider the volume of decisions made and how best to establish a proportionate monitoring system. While some models will produce decisions which can be scrutinised on an individual basis, others may need to be overseen at an aggregated level (e.g. an algorithm executing large numbers of trading orders each second). Irrespective of approach, firms will need to evidence robust frameworks for ongoing monitoring and a general understanding of the model's workings in order to satisfy stakeholders.
- Once firms have determined the appropriate level at which to explain the AI model, they will need to assess the consequences of any trade-off between explainability and the ability to generate rapid and cost-effective decisions.

2. Bias

> "In AI and ML, there is a reliance on data, but when there are biases in data or algorithms, or the situation isn't captured by past experience prediction becomes difficult and judgement becomes more important."
>
> **Mark Carney, Governor of the BoE**

Bias in AI systems arises when strategies governing the models, or the data inputs used for it, are based on unrepresentative, incomplete or incorrect criteria which can in turn drive discriminatory decision-making. This can lead to undesirable outcomes, including reputational damage, increased operating costs, service breakdown and financial loss.

AI models have the potential to promote a more objective, efficient and transparent approach to decision-making based on the available data, bypassing the outcomes that arise from human error and prejudice. However, flawed strategy, poor quality data inputs and a lack of human oversight could create the conditions for bias in AI programmes to become hardwired and scaled to a greater degree than ever before.

To date, regulatory initiatives relating to bias have tended to focus on ensuring appropriate safeguards and human oversight are in place to avoid discriminatory outcomes, especially in relation to protected characteristics.

The Centre for Data Ethics and Innovation (CDEI) intends to investigate the issue of algorithmic bias in various sectors, including financial services, through a literature review, applied technical research and public engagement workshops. This CDEI intends to report on its findings in March 2020.

Key considerations for firms

- It is important to establish controls to reduce bias at the design stage and on an ongoing basis, proportionate to the scale and complexity of the system employed, and the activities carried out. This includes ensuring diversity within design and oversight teams to help mitigate against inherent societal biases.
- Technical models should be used in combination with ongoing human oversight to review or override decisions throughout the lifecycle when required. Alert systems can help to flag unusual or unexpected activity to enhance the accuracy of model predictions.
- It is also important not to conflate 'fairness' with bias - depending on the intended purpose of the AI system, the definition of what constitutes 'fairness' can vary. Firms should instead focus on establishing a specific understanding of fairness in the context of the objectives the system has been set, and monitor outcomes against this benchmark.

3. Governance

"If firms are deploying AI and machine learning they need to ensure they have a solid understanding of the technology and the governance around it. We want to see boards asking themselves: 'what is the worst thing that can go wrong' and providing mitigations against those risks."

Chris Woolard, Executive Director, FCA

The importance of good governance is central to the UK regulators' supervisory approach and is a key consideration for them when considering firms' use of AI. Fundamentally, governance refers to the oversight and controls present within a firm. This means that every element of an AI model, from design and data training to operation and evaluation, is all governed by a similar set of standards and is subject to an appropriate level of challenge and scrutiny.

Regulators are of the view that AI governance presents some unique challenges that cannot simply be met by retrofitting existing practices. The FCA and PRA have both already flagged the need for transparent processes, adequate oversight and appropriate knowledge of AI among senior management and the board. At the global level, similar views have been given by the OECD and G20 regarding the importance of embedding skills and understanding of AI at all levels of organisations.

Key considerations for firms

- Governance processes will need to evolve alongside technology from the outset. Relying upon existing mechanisms is unlikely to fully satisfy regulators.
- Regulators will be looking for firms to acknowledge that AI is not simply an issue for the Chief Technology Officer. Silos need to be broken down to help ensure the risks, benefits and utilisation of AI are understood and promoted across the entire firm, especially by the board, thus improving operational resilience.
- AI governance frameworks should take into account the whole AI lifecycle – from design and training, to implementation and monitoring. This includes the ethical standards developers are working to, model biases and limitations, data privacy and the emergency controls in place to either intervene against, switch off or roll-back an AI model.
- Human monitoring of outcomes and processes should continue and remain aligned to the risk appetite and objectives of the firm.

4. Accountability

> "In a more automated, fast-moving world of AI/ML, boards – not just regulators – will need to consider and be on top of these issues. Firms will need to consider how to allocate individual responsibilities, including under the Senior Managers Regime."
>
> **James Proudman, Executive Director, PRA[4]**

AI accountability is the identification of who can be held responsible for a decision, action or strategy determined by an AI model.

34%

of UK FS firms assign responsibility for digital/technology /innovation to a C-suite executive[5]

But this is not as simple as it sounds. What makes AI accountability a particular challenge is the independent nature of AI design, the breadth of application throughout a business, the fact that AI may be in-house or provided by a third party (outsourced) and the reliance it has on data input from users.

The lack of clarity around accountability for AI decisions was a key theme in *a* speech made by then PRA Executive Director James Proudman, who reflected that the blurring of decision-making processes between humans and AI could lead to difficulties in identifying the root causes of problems, and by extension tracing accountability to individuals. The PRA has stated that firms will need to review how they are allocating individual responsibilities, including under the Senior Managers and Certification Regime (SM&CR).

The FCA has also called for accountable individuals under SM&CR to ensure they are able to explain and justify the use of AI systems, especially at board level. In practice, this is likely to mean board members and senior management functions (SMFs) evidencing the 'explainability' of their firm's AI and understanding its involvement within their own business units, where applicable.

Key considerations for firms

- It is clear that regulators expect to see accountability for AI enshrined at the senior level, but our research shows only 34% of UK FS firms have a C-suite individual responsible for digital/innovation/technology. This could lead to a gap between regulators' expectations and the current reality.
- Establishing accountability for AI will be a challenge. Individuals with overall responsibility for a firm's AI strategy could be a Chief Technology Officer, a Chief Data Officer or a new AI-focused role altogether. Equally, use of AI may impact a wide range of other Senior Manager Functions' responsibilities, such as the Chief Risk Function, Chief Operations Function or Compliance Function. All of these SMFs will need to show they understand and have taken reasonable steps to oversee the use of AI in their areas of responsibility.
- Firms may need to consider whether key individuals involved in the design of AI tools, such as data scientists and programmers, need to be included in the certification regime due to the potential material risks their actions could pose to consumers/markets.
- The increasing importance of AI means firms' boards will need to evidence consideration of how these developments impact their organisation, counterparts and customers and, in turn, who will be accountable for any new responsibilities this creates.

[4] James Proudman has become a Senior Advisor at the Bank of England since this speech

[5] PwC Global Fintech Report 2019

5. Data protection and security

> "Trust means citizens knowing how their data is being used, how they can control its use and where the data is going…"
>
> **Elizabeth Denham, UK Information Commissioner, ICO**

With an ever increasing societal reliance upon data, the need for adequate protection and security has never been more important. For AI systems in general and specifically common applications such as the development of ML models, the use of large datasets to train and operate is vital to their success. But just as data provides the foundations to unlock the potential of AI, it is imperative that firms use data in accordance with legal requirements and ensure it is governed robustly.

Data protection and security often work together to ensure that data (especially personal data) is used in a manner that ensures the protection of individuals' privacy rights, the robustness and safety of digital systems (such as AI-based systems) and allows organisations to meet their legal and ethical obligations to stakeholders (such as customers and employees).

Regulators have been quick to recognise the data protection and security challenges that arise with the increasing use of AI/ML systems, and have undertaken a number of important steps to address these. In the UK, the ICO for instance considers AI a priority area and has dedicated substantial resources in this field in terms of research, expert working groups and regulatory guidance, which is on the horizon. Regulators are clear in terms of their expectations of organisations developing AI, in that key data protection principles (including lawfulness, fairness, purpose and confidentiality) are even more relevant in the context of AI and ML systems. This is because the volume of personal data used in their development poses a significant potential harm if misused.

Key considerations for firms

- Data protection and security risks must be assessed up-front, with the involvement of all relevant stakeholders. This includes not just the development teams but data protection, security, risk and governance specialists.
- Organisations must carefully consider the lawful basis for the use of personal data in AI/ML systems and have corresponding justifications for the basis used. Personal data collected for a particular processing activity must not be repurposed for use in training ML systems for example, where the original lawful basis does not account for such use.
- Firms must ensure that individuals can exercise their data protection rights, such as the right to be forgotten, rectification of data and human intervention/review. This includes considerations such as retention periods for training data.
- Consent for data processing may need to be multi-tiered given the number of ways customer data may be used. In addition, training algorithms and modelling outcomes may not be a sufficiently legitimate reason for processing personal data under GDPR, in some cases.
- Senior leadership should recognise the reputational, financial and regulatory impacts of an AI security breach on their businesses. For example, should an ML model be hacked and given false training data, malicious outcomes could be realised on a systemic level. It is critical to know how to prepare for an AI cyber incident, to understand where it may manifest and how to respond quickly.

What are the next steps?

Regulators in the UK and at the international level have started to set out some clear expectations around the use of AI and firms need to ensure these are factored into their AI strategies. But the debate on how the regulatory regime needs to adapt further is ongoing. The BoE and FCA have announced they will establish a public-private forum to determine an appropriate supervisory approach to AI and ML and John Glenn, City Minister at HMT, has stated that the UK's regulatory regime will need to respond to developments in AI[6], so this is an area which is likely to continue evolving in the UK. We also expect to see increased output from the international standards setters in 2020, so firms need to follow these developments closely.

For firms, given that the regulators' approach is largely exploratory in nature at this stage, there is a clear opportunity to step up and take the initiative individually or collectively with peers. Regulators will be receptive to collaborating with those closer to the details of implementing AI systems to help broaden their understanding of the potential implications and help create proportional standards for the use of AI.

How can PwC help?

At PwC we believe a holistic approach is key to successfully designing, building and embedding AI and ML into your business. We have developed a comprehensive package of support which helps firms get AI projects off the ground and deploy them at scale in a responsible way.

- **Bias and fairness** – We can help organisations to define, detect and intervene against bias, taking corrective action to improve the decision-making process.
- **Data protection and security** – Support can be provided to identify areas of weakness in models, assessing system safety and establishing systems that produce reliable, consistent outcomes over time. We also provide tools to improve GDPR compliance and cyber security in firms, including simulating security breaches for senior management in our interactive Game of Threats.
- **Explainability** – We can help firms to explain and defend business-critical decisions and individual predictions, tailored to the requirements of specific stakeholders. Our Explainable AI framework includes use case criticality, which can support you in identifying the impact of every wrong prediction in the context of regulatory requirements, reputation, risk, rigour and revenue.
- **Governance and accountability** – Our Responsible AI Framework helps firms to establish an enterprise-wide governance and accountability framework for AI applications throughout their lifecycle, ensuring consistent and proportionate systems and controls are in place to identify and mitigate risk. This includes integration with SM&CR requirements. We can also provide advice and/or challenge to help you agree a set of principles for the compliant and responsible adoption of AI in your firm. Within our 'Ethical AI' framework, we have conducted a market scan and academic research on existing codes of conduct and AI principles which can be tailored to define a set of principles for you.
- **Horizon scanning** – We can identify and evaluate relevant regulations, regulatory guidance and ethical considerations that impact AI solutions in your business – allowing you to start on the right foot and reduce regulatory risk.

[6] https://www.gov.uk/government/speeches/the-uk-financial-services-beyond-brexit-summit-economic-secretarys-speech

Contact details

Leigh Bates
Partner
T: +44 (0)7711 562381
E: leigh.bates@pwc.com

Tom Boydell
Manager
T: +44 (0)7483 399332
E: tom.boydell@pwc.com

Conor Macmanus
Director
T: +44 (0)7718 979428
E: conor.macmanus@pwc.com

Leo Donnachie
Senior Associate
T: +44 (0)7483 329595
E: emilio.donnachie@pwc.com

191125-090456-TB-OS