Minutes
Artificial Intelligence Public-Private Forum - Second meeting
26 February 2021

Attendees

Co-Chair	Organisation
Ramsden, Dave	Bank of England
Mills, Sheldon	Financial Conduct Authority

Moderator	Organisation
Cook, Nick	Financial Conduct Authority

Member	Organisation
Baldwin, Michael	Google Cloud
Barto, Jason	Amazon Web Services
Browne, Fiona	Datactics
Campos-Zabala, Javier	Experian
Christensen, Hugh	Amazon Web Services
Dorobantu, Cosmina	Alan Turing Institute
Gadd, Sarah	Credit Suisse
Kellett, Dan	Capital One UK
Kirkham, Rachel	Mindbridge AI
Kundu, Shameek	Truera
Lennard, Jessica	Visa
Moniz, Andy	Putnam Investments
Morris, Owen	Aviva PLC
Morrison, Gwilym	Royal London
Rees, Harriet	Starling Bank
Rosenshine, Kate	Microsoft UK
Sandhu, Jas	Royal Bank of Canada
Shi-Nash, Amy	HSBC
Tetlow, Phil	IBM UK
Treleaven, Philip	University College London

Observer	Organisation
Beswick, Jacob	Office for Artificial Intelligence
Buckley, Oliver	Centre for Data Ethics and Innovation
Dipple-Johnstone, James	Information Commissioner's Office
Mountford, Laura	HM Treasury
Yallop, Mark	FICC Markets Standards Board

As stated in the AIPPF Terms of Reference, the views expressed by the AIPPF members in these minutes and all other outputs do not reflect the views of their institutions, the Bank or FCA. Also, the activities, discussions, and outputs of the AIPPF should not be taken as an indication of future policy by the Bank or FCA.

Item 1.   Opening remarks by co-chairs and ICO Observer

Co-chairs Sheldon Mills and Dave Ramsden welcomed the members and observers to the second meeting of the Artificial Intelligence Public-Private Forum (AIPPF), before James Dipple-Johnstone of the ICO gave some introductory remarks.

Sheldon Mills

Sheldon noted that data, the main subject of this meeting, is at the core of financial services. This is true of the past and, increasingly, for the future of financial services. Artificial intelligence (AI) and machine learning (ML) are accelerators of this broad trend. Consequently, the Financial Conduct Authority (the FCA) and the Bank of England (the Bank) are working to understand the changing role of data and data-driven technology – not just in terms of what it means for markets but also in harnessing the benefits to become more effective, data-enabled regulators.

Sheldon welcomed the publication of [the Kalifa Review of UK fintech](#) as an important contribution in helping to shape and build on the UK's successes to date in financial services innovation. He highlighted the AIPPF as a way of harnessing the expertise of a broad range of stakeholders and using collaboration to identify key challenges and develop responses.

Sheldon welcomed James Dipple-Johnstone, Deputy Commissioner at the Information Commissioner's Office (ICO) and AIPPF observer, to this discussion, noting that data protection is at the heart of the data economy.

Dave Ramsden

Dave reiterated Sheldon's comments on the importance of data for financial services and went on to outline why the topic of data is important for the work of the AIPPF, the Bank and, ultimately, the safe adoption of AI in UK financial services.

On the importance of data for the AIPPF, Dave reiterated that the aim was to get to the heart of the key data-related risks and challenges related to the use of AI in financial services. However, the purpose of the AIPPF is also to move the dialogue forward and explore potential solutions to these problems.

As for the importance of data for the Bank, Dave explained that the Bank and FCA had launched a consultation on the transformation of data collection over the next decade. The recent Bank and FCA [letter to chief executive officers](#) provided an update on that work. The Bank is also undergoing a process of making the [PRA's Rulebook machine-readable](#) and using ML tools to enhance the capabilities of supervisors.

In addition to the internal data-related work, the Bank and other regulators should consider how best to support firms' safe adoption of AI and what that means for the relevant regulatory frameworks. Although the AIPPF is not a decision-making body and the outputs should not be seen as an indication of future policy, part of the AIPPF's purpose is to understand how existing policy frameworks affect and encompass AI, and what the appropriate level of any potential policy should be.

<u>James Dipple-Johnstone</u>

James noted that, as the UK data protection authority, the ICO's interest lies in both understanding and supporting confidence in the use of personal data. The ICO recognises the potential of new technologies, such as AI, and the benefits they can bring. But it also recognises that public confidence is key to the adoption of AI and other new technologies.

James said the ICO is also considering how it can develop its own services to meet its joint statutory objectives of upholding rights and supporting economic growth. The ICO has been building up its capabilities in the AI space and working on its regulatory approach. This includes consulting on how auditing of AI applications may work in practice, so the ICO can better evaluate compliance cases as well as provide advice and a [safe space for firms to test things out](#).

Lastly, James reiterated that the ICO is working closely with other regulators both domestically and internationally to ensure there is a consistent and coordinated approach, and that it develops initiatives in an appropriate way.

Item 2. Roundtable discussion

The aim of the roundtable discussion was to identify and discuss the key issues and challenges for each of four topic areas:

1. Data quality
2. Data strategy and economics
3. Data governance, ethics, and culture
4. Data standards and regulation

1. Data quality

Key challenges

1.1. Issues arising from the use of 'alternative data' (including unstructured, synthetic, aggregated, and third-party data).
  - The amount and associated costs of data cleansing needed for effective use of unstructured data.
  - Creating synthetic data that capture underlying structures, correlations, and relationships of real data.
  - Validating aggregated data without knowing the granular structure of those data
  - Ensuring the quality, provenance and legality of third-party data.

1.2. Adapting existing data quality standards to an AI context including management of data provenance and lineage.
  - Applying existing data standards to AI contexts.
  - Lack of an industry-wide consensus on data standards in financial services, including agreement on good practice.
  - Ensuring better understanding of data provenance and lineage at board level.
  - Key roles required and the importance of data stewardship to ensure data standards remain high.

1.3. Lifecycle management including data aggregation, transformation, and integration of diverse sources (e.g. public vs. private data).

- Feature engineering is crucial and at the heart of AI development.
- Ensuring that appropriate and effective processes are in place for building, using, maintaining, and documenting features.
- This is one area among others where there may be a lack of awareness and misunderstanding.
- Potential for multiple similar features being constructed, which in turn may lead to confusion and errors.

1.4. Versioning, documentation, inventory, and reporting challenges specific to AI systems.

- Ensuring that an appropriate and effective model inventory management process is in place to provide a clear understanding of where AI is used within organisations.
- Importance of documentation and versioning for models, code, and data.
- Lack of appropriate data science skills can also be a challenge.
- Ensuring consistent monitoring and understanding of how an AI model may change over time, in particular when presented with new data.

Discussion

1.5. Members debated the root causes of the above data quality challenges and if these related to knowledge of what is required, skills and resourcing, or sponsorship and ownership. They noted it was generally a combination of all three but the root causes can differ depending on the maturity of an organisation's use of AI. There has been good progress on addressing the skills gap but there are still issues when applying data science and AI skills across business lines.

1.6. In terms of sponsorship and ownership, ensuring board level appreciation of data quality issues is always key. However, this can be an area of confusion especially when taking a retrospective rather than a proactive approach to data quality. Noting that non-AI models have historically been used and governed in relatively narrow ways within financial services firms, the step-change with AI means data quality management and ownership should be important considerations for corporate decision-making from the outset.

1.7. Members noted that, in general, a key driver of many data quality challenges can be an attempt to retrofit existing process, controls and systems to AI. Pre-existing ways of managing and governing data quality may not necessarily be appropriate for AI models. This can potentially create a problem where existing control frameworks cannot scale to the volume and variety of features in the data or the range of applications of the AI model. Scaling may also lead to changes in data quality attributes, especially when using narrower definitions of data quality from pre-existing approaches. For example, the increased relevance of data completeness and ensuring data is representative were noted as important considerations for AI models, above and beyond well-known data quality matters such as accuracy and timeliness. The Alternative Data Council has done work to establish best practice and data quality standards for alternative data, which may be useful for financial services firms.

1.8. Unlike traditional models, AI models may take less time to develop and put into production but may need continuous monitoring. For example, accurate version control and documenting changes for unsupervised learning models that are constantly changing would require different approaches. Should the focus be less on version control and more on outcome metrics to facilitate real-time monitoring of performance? Similarly, there may be a need for systems-level testing for more complex models. Members suggested that a mental shift is likely needed to move from the more traditional waterfall approach of model development and production to working with AI models.

1.9. The importance of data quality is likely to increase with further data sharing and open finance initiatives. High-quality data often come from sources that have an interest in curating their data but that may not necessarily be the case for all organisations. In effect, achieving high-quality of data is not cost-free and some approaches to open data sharing may disincentive entities from investing sufficiently in assuring the quality of their data assets. Data quality can therefore differ significantly, in particular in an open data environment. Work being done on data standards as part of [Open Banking](#) may be useful when considering AI in financial services.

2. Data strategy and economics

Key challenges

2.1. Data value, data economics, and business model considerations including aligning data strategy to business strategy.
- Fair and equal access to data.
- Pricing of data.
- Questions of consistency and transparency regarding third-party data.
- Estimating expected returns from data, given associated costs.

2.2. Decentralised and networked data systems, such as the role of federated learning and bringing computation to data.
- Some network nodes become important purely because of their position in the network rather than any innate properties.
- Latency of information retrieval and transaction speeds across large networks.
- Emergent behaviour (when a number of simple entities operate in an environment and form more complex behaviours) in networks both at micro and macro levels.
- Emergent behaviour is very difficult to predict and its implications difficult to project and analyse.

2.3. Intersection of data management, model management, and risk management (including the challenges of working and sharing data across multiple jurisdictions).
- Data quality management, including ownership and accountability.
- Managing issues like look-ahead bias and survivorship bias.
- More clarity may be needed on what constitutes 'inside information' in the context of alternative data.
- Differences in incentive structures (e.g. between finance professionals and data scientists) across organisations may result in the implementation of different practices.

2.4. Measuring and tracking data flows within as well as in and out of the organisation including real-time and streaming data collection.

- Implementing and governing liability, control, and ownership of data in a highly complex distributed system where data are constantly transferred and modified.
- Move towards edge computing may have consequences for firms and regulators outside core banking platforms.
- Data audit, both in terms of who (human vs machine) and when.
- Data use – how much has been used and does usage align with expenditure?

Discussion

2.5. Members commented on the challenges associated with data streaming and real-time analytics, which are largely related to the underlying infrastructure and technology stack. However, some firms are able to stream customer data from various sources, link it together and make the data available to business users so they can act on the intelligence immediately. Such processes can be automated and may provide more efficient customer propositions.

2.6. There are challenges in accurately assessing the overall costs (and expected benefit/value-add) of individual data assets. These challenges become even greater if data marketplaces include monopolistic or restrictive tendencies, or where, as a result of market structure, access to datasets may be limited. This may raise challenges for pricing policies.

2.7. There are also challenges associated with understanding the provenance and legal status of data sourced from vendors that scrape website data or collate it from a range of sources. This can create risks for firms and may have implications for consumers including questions on what data customers are willing to give up for free (e.g. via social media sites or consent to all cookies), how that may be used both in financial services and other sectors, and the concentration of those data.

2.8. Members also discussed the need to create a framework to audit AI. This was seen as achieving assurance within the firm but also as a way to contribute to building trust in the technology more broadly. Members noted that there is currently no best practice available on AI auditing, and that there are many outstanding questions. How holistic should the auditing framework be? Should it focus on the model or the data (or both)? What should it audit for (including but not limited to accuracy, explainability or fairness)? And whose responsibility is it to develop this ?Members also noted the challenge of aligning AI auditing models with other applicable standards – such as sectoral guidance, data protection and international alignment of approaches.

3. Data governance, ethics and culture

Key challenges

3.1. Approaches to measuring and quantifying ethical dimensions, including the challenges of operationalising ethical principles.

- AI algorithms span a wide range of methods and use-cases; accordingly, the level of ethical oversight could differ depending on the use-case and associated risks and benefits.

- Human inputs are vital in understanding statistical outputs and querying the underlying data, such as why certain data were sampled.
- Getting a diverse team and set of views to review algorithms and data from an ethical standpoint.
- Operationalising and measuring the implementation of AI ethics, which requires education of the principles at every level of the organisation and very clear definitions at every level of the process.

3.2. Data ownership and accountability structures across the organisation, including consent management.
- Split between data ownership and AI ownership. Should firms create data roles, AI data roles, or joint roles?
- Lack of access to data, which creates a risk of inconsistent behaviours and processes as the quality of the data may vary, thereby potentially introducing statistical bias.
- Data silos within organisations.
- The need for human-in-the-loop mechanisms.

3.3. Managing data privacy, access, and availability
- Ensuring full compliance with data protection and additional requirements related to the use of customer data.
- Privacy risks that may emerge when individual datasets are combined with other (internal or external) sources that are not covered by the same decision process.
- Basic data literacy to empower employees to both understand data and how to use it, but also to understand how it is used.
- Bottom up and top down approaches to data management.

Discussion

3.4. The discussion started by highlighting several aspects of the European Commission's 'Ethics guidelines for Trustworthy AI' that are relevant to the AIPPF and the topic of data. The first is the need to ensure human oversight or a human-in-the-loop as part of data and AI governance. Second, it is important that AI is not represented as being human to the customers. Third, AI algorithms should be considered as stakeholders for other AI algorithms.

3.5. Members noted that conflation of the terms 'fairness' and 'bias' can be a challenge. There are numerous definitions of fairness; which one to use is often a business decision reflecting the goals and outcomes of the business. In addition, businesses may be uncomfortable discussing fair outcomes and may not have a clear philosophy on outcomes, compared to fair processes.

3.6. AI is pushing up against this issue, which may see institutions having to consider fair outcomes and take value-based decisions. AI may impel institutions and users to define fairness in mathematical terms and code it into their models, in so far as that is possible. There has already been much work on treating customer fairly in financial services, and it may be possible to expand and build on it to incorporate AI.

3.7. Change management in organisations is a non-trivial aspect for fairness and bias considerations. Specifically, creating the right environment for conversations on ethics takes time and requires buy-in from senior leadership. The right skillsets are needed to

discuss ethics and fairness and a multi-disciplinary approach is needed. Again, the use of third-parties can create additional challenges, since vendors and customers can have different ways of assessing fairness. Similarly, what may be considered fair can change as society and demographics change over time.

3.8. A further question is whether fairness and bias apply to businesses and corporate entities, rather than just to individuals. The forum discussed whether obligations of fair treatment by corporate institutions applied to their interactions with corporate clients, or only to retail clients. The group acknowledged that within certain legislation and regulation, small and micro-businesses have similar rights to natural persons. How concepts of fairness apply beyond natural persons could, for example, have implications for lending to small businesses versus large businesses.

4. Data standards and regulation

Key challenges

4.1. Development and use of data standards.
   - Incremental standards may be needed for most elements of traditional data governance.
   - These include: data documentation, quality, retention, privacy, sovereignty, and security.
   - In addition to these data standards, there are challenges around the adoption of standards to enable development and roll out of technologies needed to make data available for AI models.
   - These include: differential privacy, federated learning, homomorphic encryption, personal data wallets and federated consent management. Most of these are not unique to AI but their importance increases significantly in the AI context.

4.2. Equal access and fair pricing of third-party data.
   - How to achieve pricing of third-party data based on the value those data generate.
   - Would a push towards equal access to third-party data disrupt the potential for market forces to identify high quality datasets and providers, and filter out low quality ones.
   - Potential for equal access to third-party data could cause incremental risk in data handling as data is transferred between multiple parties in the value chain.
   - Ensuring data regulators and financial service regulators stay aligned.

4.3. Data auditing, certification and attestation.
   - Lack of certification for data sets, including for provenance, ageing, how data are generated, what changes were made to the lifecycle and what the known limitations are.
   - With respect to auditing, there is a need to consider data and model auditing holistically.
   - How to fit any new AI auditing framework into an organisation's existing Enterprise-Wide Risk Management frameworks.
   - On certification, there are challenges around who the certification would be aimed at, who does the certification, against what standards and how to prevent the loss of Intellectual Property through certification.

4.4. Where additional guidance on existing regulation (both UK and international) could be useful, the appropriate level of regulation and what should be included in the scope.
- At the international level, there is value in multilateral alignment for both data and AI governance.
- Industry-agnostic coordination but also coordination across financial services regulators both at the domestic and international levels.
- Important for data privacy regulators and financial service regulators to work together in creating a coherent sectoral set of standards.
- Any guidance should not be overly prescriptive and should provide case studies to illustrate what 'good looks like'.

4.5. Creation and use of cross-organisational datasets that can support the public good and increase financial inclusion.
- Cooperating on 'data for good' is extremely valuable, as underlined by work on the pandemic response.
- But it can be difficult to achieve for a range of reasons: privacy (including consent), commercial considerations, disclosure requirements, interoperability, data quality and security.

Discussion

4.6. One of the key questions for the AIPPF is 'what changes with AI' with respect to the broader discussion on data standards, risk management and auditing.

4.7. There is a difference between regulating models themselves with a high-level of explainability versus treating them as black boxes and regulating their inputs, outputs and outcomes. Which approach is more appropriate depends on the context, use-case, and materiality of the model. Similarly, there are certain jurisdictions where the regulations are more prescriptive for models and those guidelines will require higher levels of explainability. Having clear guidelines could increase confidence when deploying the technology, but could equally hamper desirable innovation. Striking the right balance is a key consideration for regulators and policy makers.

4.8. The ICO guidance on explainability may be an example that strikes the right balance on levels of explainability, depending for example on the use case and the individual seeking the explanation. Ultimately, the aim is to provide reassurance around the governance and provenance of data, and the link between input and output, without having to delve into the black box.

4.9. Auditing may build trust and confidence in AI. In order to ensure an effective AI audit, a framework must be in place around which assurance products could be developed. There are also lessons to be learned from financial audit and control frameworks, including good practice for controls and the use of data and AI.

4.10. There may also be space to create a risk management and governance framework for AI-related issues. One approach could be to develop a set of AI risk principles and map them against firms' existing risk frameworks. This would enable firms to identify any potential risks that aren't already covered and focus efforts on those issues within the AI governance framework. This approach could also help focus on areas where staff needed to be upskilled, such as those in the model risk management teams.

Item 3.   **Closing remarks and next steps from the moderator and co-chairs**

The moderator and co-chairs thanked the members and observers for the engaging conversation as well as their continued support. He outlined the next steps, which include preparations for the forthcoming workshops on data in March and the subsequent AIPPF meeting on model risk management in Q2.