



College of Humanities
Master of Financial
Technology-
Fin Tech

Artificial Intelligence for Financial Technology

FNT 626

Lecture 03

Prepared By:
Dr. Abeir Moussa



Artificial Intelligence for Financial Technology

FNT 626- AI for Fin Tech

Lecture 03



Responsible AI: An opportunity to test regulatory principles

LECTURE 03

Ch.4- Responsible AI: An opportunity to test regulatory principles

- **Learning Objectives**

1. **Understanding Responsible AI Principles**
2. **Exploring Regulatory Frameworks for AI**
3. **Examine the implications of AI on traditional regulatory principles, such as data privacy, consumer protection,**
4. **Explore the ethical and Social Considerations in AI**
5. **Challenges in Implementing Responsible AI**
6. **Analyze real-world case studies and best practices in responsible AI implementation,**
7. **Explore the future Trends in Responsible AI Regulation**

Ch.4- Responsible AI: An opportunity to test regulatory principles

- **Chapter outline**

1.INTRODUCTION,

2. FROM PRINCIPLES TO PRACTICE: AN INITIATIVE TO EXPLORE THE APPLICATION OF REGULATORY GUIDANCE ON RESPONSIBLE AI TO FINANCIAL SERVICES USE CASES,

3. LESSONS LEARNED FROM COLLABORATIVE INITIATIVES ON THE APPLICATION OF RESPONSIBLE AI PRINCIPLES,

4. CONCLUSION

Chapter Outline

1. INTRODUCTION

1.1 Regulatory Principles as a First Step

1.2 Dialogue Among Industry and Regulators



LECTURE 03

Ch. 4- Responsible AI: An opportunity to test regulatory principles

Ch. 4- Responsible AI: An opportunity to test regulatory principles

INTRODUCTION

There is strong interest from the financial services industry in deploying AI to aid financial institutions' (FIs) digital transformation, a competitive imperative today. As of April 2021, more than 60% of financial services organizations had begun to accelerate their pace of digitalization in response to the pandemic. Spending on AI by the financial services sector in the region has significantly risen and is estimated to reach USD 4.9 billion by 2024. Currently, AI is primarily being used in the financial services industry in three broad categories: customer experience, employee empowerment, and AI-empowered insights. Those FIs that have adopted AI across different business functions have seen improvements of between 17% and 25% in areas such as better customer engagement, higher competitiveness, accelerated innovation, higher margins, and improved business intelligence.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

1.1 Regulatory Principles as a First Step

Corresponding to this growth of AI in the financial services industry, there has been a surge in public discussion on how to responsibly implement these technologies. Regulators have generally proceeded deliberately in this area.

One of the first pieces of guidance on responsible AI was the Monetary Authority of Singapore's (MAS's) Principles to Promote **Fairness, Ethics, Accountability and Transparency** in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector (FEAT Principles) in 2018.

MAS explained that the FEAT Principles were 'not intended to be prescriptive', and 'through the process of ongoing industry engagement, there is scope for MAS to identify areas where more specific or technical guidance would benefit the industry, with this set of Principles serving as a foundational framework for consideration and discussion'.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

The **Hong Kong Monetary Authority (HKMA)** released its High-Level Principles on Artificial Intelligence in 2019, as well as separate guidance to increase FIs' understanding of consumer protection aspects of AI applications. Similar to MAS, the HKMA explicitly stated that 'these principles are high-level in nature as the HKMA is mindful that overly prescriptive or rigid requirements may inhibit the further development of AI-related technologies.' Financial services regulators in Germany, the Netherlands, and the US have developed similar principles addressing the use of AI in financial services, as opposed to mandatory regulatory requirements. The proposed regulation seeks to provide AI developers, deployers, and users with clear requirements and obligations regarding specific uses of AI.

It was not until 2021 that the European Commission proposed the world's first mandatory regulation setting forth rules on AI across all industries. The explanatory memorandum describes how the proposal draws on a considerable amount of prior work, including the European Commission's February 2020 white paper on AI, which was followed by a 'broad stakeholder consultation'. The Commission also noted that the European Parliament adopted several resolutions related to AI in October 2020, including on ethics and AI, which informed the proposal.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

1.2 Dialogue Among Industry and Regulators

The principles-first approach taken by regulators has opened a forum for collaboration between regulators and industry – with that dialogue allowing regulators to refine and converge on basic principles of responsible AI, as well as gain an insight on whether new and specific regulation is needed to govern the use of AI.

MAS's Veritas Consortium is an example of this, with MAS actively engaging with industry participants such as SGInnovate, EY and financial institutions to develop assessment methodologies for the principles in the FEAT Guide, culminating in the publication of white papers in February 2022. The feedback MAS received throughout the Veritas process will provide it with important learnings for any future guidance. The FEAT Guide itself was 'co-created with the financial industry and other relevant stakeholders', in close collaboration with the Singapore **Personal Data Protection Commission (PDPC) and Infocomm Media Development Authority (IMDA).**

Ch. 4- Responsible AI: An opportunity to test regulatory principles

In the United States, leading organizations across the financial services, technology and academic industries announced the formation of a new **National Council for Artificial Intelligence (NCAI)** that will look to co-create AI solutions, identify and set the AI strategy and vision for a wide range of projects and track progress towards AI adoption.

There also seems to be increased coordination among different jurisdictions regionally and globally, for example, through the **Basel Committee** and Global Partnership on AI. At least in part due to this ongoing dialogue among regulators and industry across borders, today there is a high level of convergence on the underlying principles and framework for governing the responsible use of AI. It makes sense that at this nascent stage of AI use, governments have proceeded deliberately in refraining from issuing new regulations specific to AI, instead working first to establish non-binding principles. This principles-first approach, combined with ongoing dialogue among various regulators and industry and technology providers, has allowed regulators to refine and converge on basic principles of responsible AI which can serve as a foundation for any future regulation.

2. FROM PRINCIPLES TO PRACTICE: AN INITIATIVE TO EXPLORE THE APPLICATION OF REGULATORY GUIDANCE ON RESPONSIBLE AI TO FINANCIAL SERVICES USE CASES

2.1 Use Case 1: Using AI to Predict Which Cards Will Be Used for Travel

- 2.1.1 Assessing the materiality of the use case
- 2.1.2 The ongoing relevance of existing internal risk governance and regulatory frameworks
- 2.1.3 Ensuring fairness is achieved in the context of the use case
- 2.1.4 The collaborative role between financial institutions and technology partners to promote transparency.

LECTURE 03



Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2 Use Case 2: **Using AI to Automate Verification of Wet Ink Signatures**

2.2.1 Assessing the materiality of the use case

2.2.2 The importance of context in determining which principles are relevant

2.2.3 The balancing act between promoting transparency and ensuring the security of AI systems\

2.2.4 Guardrails in place to promote fairness in the AI system

2.3 **Use Case 3: Using AI for Regulatory Compliance Relating to Know Your Customer (KYC) Checks**

2.3.1 Assessing the materiality of the use case

2.3.2 The ongoing relevance of existing regulatory guidance and frameworks

2.3.3 Fairness and how biases can be exasperated and mitigated

2.3.4 Potential trade-offs between transparency and the security of AI applications



LECTURE 03

Ch. 4- Responsible AI: An opportunity to test regulatory principles

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.FROM PRINCIPLES TO PRACTICE: AN INITIATIVE TO EXPLORE THE APPLICATION OF REGULATORY GUIDANCE ON RESPONSIBLE AI TO FINANCIAL SERVICES USE CASES

Using this opportunity, Microsoft spearheaded an initiative in collaboration with Standard Chartered, Deutsche Bank, Visa and Linklaters to explore the application of the MAS FEAT Principles and the PDPC AI Framework (together, ‘responsible use principles’) to three real-world AI use cases in the financial services industry. For each use case, one financial services partner led the assessment of how the responsible AI principles would apply to that specific adoption of AI technology, and the whole group then discussed and debated the findings. The aim of the group of partners was not to identify a definitive set of methods for implementing responsible AI, but instead to provide real-world insights into the benefits, issues and challenges identified in the implementation of AI principles.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.1 Use Case 1: Using AI to Predict Which Cards Will Be Used for Travel

The first use case involved the use of AI to predict which bank cards would be used for upcoming travel. For this use case, an organization was applying AI to its deep dataset of card transactional data to be able to predict which cards are likely to be used for travel in the coming months. Under a business-to-business (B2B) service, the organization provided that insight to the card issuers, who matched the travel propensity scores with their own datasets to determine the ideal campaign recipients of travel-related marketing. Importantly, no protected attributes such as personally identifiable information like names, gender, or age are used in the AI model-building process for this solution.

The use case raised interesting discussions around the materiality of the AI system, as well as the applicability of internal risk governance and regulatory frameworks in an AI context, the relevance of certain responsible AI principles compared with others, and the guardrails needed to be implemented to ensure the system was fair and transparent in nature.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.1.1 Assessing the materiality of the use case

The analysis began with a materiality assessment to determine the proportionate application of responsible AI principles. Considerations that formed part of this materiality assessment were how the results of the model are being used, and what data is used to train the model. It is also important to foster an organizational culture and account for diverse viewpoints in these risk assessment decisions to ensure that responsible use principles are appropriately applied, and that potential concerns are identified and dealt with early.

Overall, it was observed that any incorrect decision made by the AI application in this case would have a low probability and severity of harm to that individual. This determination means this use case would not require an intense review as compared with a more potentially harmful AI application.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.1.2 The ongoing relevance of existing internal risk governance and regulatory frameworks

Another consideration raised in this use case was whether any of the organization's internal risk governance procedures were relevant. The organization had an internal governance framework that dictates clients cannot be given information that will grant a client a competitive advantage over another or allow a client to understand how its competitors are faring in the market.

These internal governance considerations led to an analysis of the type, quality and quantity of the datasets being fed into the AI model. This in turn addressed some of the expectations found in the MAS FEAT Principles and the PDPC AI Framework around the Accountability and Ethics principles. So, this example highlights how internal governance considerations remain relevant in ensuring that AI development aligns with the organization's ethical values and principles. Another important consideration is whether any existing regulatory frameworks are relevant in the context of the AI use case.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

In this case, data protection regulations prescribe a legislative requirement for the issuing banks to obtain the necessary consents from their customers for data processing activities, and before tailored marketing and advertising campaigns can be sent to them.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.1.3 Ensuring fairness is achieved in the context of the use case

One of the responsible AI principles that was relevant in this use case was fairness. Like other principles, fairness is not only dependent on differences in context, but also in time. In the context of this use case, travel predictions for the issuing bank's customers may vary from time to time and are subject to seasonality. There is a danger that individuals will be disadvantaged by decisions made in a particular season by the AI solution which is trained on data gathered from another season. To ensure fairness in decision-making, the model should account for seasonality and be refreshed periodically to ensure new trends are captured to keep predictions accurate.

Fairness considerations may also raise a dichotomy between group and individual fairness. In this use case, group fairness would involve sending offers to demographic groups that typically do not receive them. Individual fairness would, however, entail only sending offers to specific customers that the model predicts are likely to travel. Consideration needs to be given before and during deployment of the AI system to assess the appropriate balance.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.1.4 The collaborative role between financial institutions and technology partners to promote transparency

Another responsible AI principle that arose in this use case was transparency – and how the relevant stakeholders could meet their transparency obligations through collaborative efforts. A B2B service would particularly benefit from collaborative efforts with technology partners to promote transparency. As a B2B service, the organization has developed and applied an AI system that makes decisions affecting end consumers, but will not itself have access to, or interface with, the end consumer.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

It was determined the best approach would be for the AI developer to provide technical information – such as model predictors and fully documented model inputs and outputs. This allows the AI deployer to then educate end users on how decisions are made using AI. The AI deployers are also empowered to check model quality and provide feedback on model performance to the organization for the purposes of improving and refreshing the model. In addition to this technical information, the AI developer could make available non-technical documentation to explain how the system works, so that the AI deployers could use this in explaining the functioning of the AI system to consumers in plain and understandable terms. Collaborative efforts between AI developers and deployers are therefore crucial to ensure each stakeholder fulfils their transparency obligations

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2 Use Case 2: Using AI to Automate Verification of Wet Ink Signatures

The second use case involved the use of AI to automate the verification of wet ink signatures. In this use case, an organization has developed an AI solution that aims to enhance an existing manual operational process – the verification of authenticity of wet ink signatures. Key issues that surfaced, and that are discussed below, were the materiality of the use case, the importance of context in evaluating regulatory expectations, the applicability of certain principles in the responsible AI guidance, and what requirements would be needed to promote fairness and transparency upon deployment of the AI system.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2.1 Assessing the materiality of the use case

A materiality assessment of the use case was first conducted to determine the proportionate application of responsible AI principles. The probability and severity of harm to an individual because of the use of AI was determined to be on the lower end of the spectrum. The AI system was being used to support or replace an operational decision as to whether a client instruction should be processed or escalated to a human for review, and so the risk of harm to the individuals concerned seemed to be minimal.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2.2 The importance of context in determining which principles are relevant

In this use case, there were certain measures that the MAS FEAT Principles and the PDPC AI Framework guidelines suggested organizations should implement which did not seem relevant in the context of this particular use case. For instance, the principle to provide consumers with an explanation on what data is used to make artificial intelligence driven agent (AIDA) decisions about the consumers is not relevant in this use case given that the only input data is their wet ink signatures. Similarly, the principles concerning the use of personal attributes of data subjects in ways that may generate bias were not relevant as there were no personal attributes being used for this particular AI application, beyond the wet ink signature itself. This underlines the importance of context in determining which principles are relevant in a specific use case.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2.3 The balancing act between promoting transparency and ensuring the security of AI systems

One of the responsible AI principles that was relevant in this use case was the transparency principle. The broad objective of the transparency principle is to increase public confidence and trust in the use of AI. To enact this, the organization would need to ensure frontline staff are equipped to explain the use of AI to customers in a simple yet adequate manner. An important area of discussion was the level of detail that needed to be provided to customers to prevent any 'gaming of the system' from information disclosure. Again, the context of this use case was relevant and meant that while some features of the transparency principle was applicable (as discussed above), other aspects were not. For example, the use case did not warrant a substantial application of the accountability measures that are set out in the regulator guidance, such as the need for alternative channels for redress.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

2.2.4 Guardrails in place to promote fairness in the AI system

Fairness was also found to be relevant in this use case. The AI application was found to enhance fairness in itself – as using AI instead of a manual, human-driven process to compare signatures reduced some of the disadvantages experienced by aging or disabled customers. This is because the AI system is designed to detect changes in signature patterns over time, and prompt customers when the time has come for a specimen signature on the system to be updated.

Nonetheless, it was important to consider if any other guardrails were needed to ensure the AI application was fair. It was decided a human-in-the loop approach would be useful for this purpose. This would involve the incorporation of a differential tolerance (i.e., how closely matched the signatures are) so matches that fail below this threshold are escalated appropriately.

2.3 Use Case 3: Using AI for Regulatory Compliance Relating to Know Your Customer (KYC) Checks

The third use case involved using AI in know your customer (KYC) checks to promote regulatory compliance. To build a trusted KYC process, financial institutions (FIs) need to capture quality and complete customer data in the shortest possible time. The AI KYC model in this use case involved both internally and externally sourced datasets to form and train the models.

Key considerations that arose, and that are discussed in detail below, include the materiality of the use case, the relevance of existing regulatory guidance or frameworks to an AI context, the partial relevance of principles found in the responsible AI guidance, and the guardrails needed to promote fairness and transparency in this AI system.

2.3.1 Assessing the materiality of the use case

A key consideration was the materiality of the use case, and how that would dictate the proportionate application of responsible AI principles. It was determined that the materiality of the decisions made about data subjects in this case was high. This was because the AI model is determinative of whether financial products/services are offered to the data subject (i.e., whether a loan or line of credit is approved).

The source of input data is also a relevant factor in determining materiality. A tension was raised between two considerations: on the one hand, organizations needed as complete a dataset as possible to produce accurate results; while on the other hand, using more datasets from third parties can increase the risks of poor dataset quality, the potential for false positives, and unintended bias in the results.

LECTURE 03

Ch. 4- Responsible AI: An opportunity to test regulatory principles

Gathering more data may also go against the privacy principle that personal data collection should not be excessive. This was deemed to raise the gross materiality assessment, as there is a significant risk of harm if appropriate due diligence procedures and internal risk assessments to track data lineage and to verify data sources are not implemented. While this risk would similarly exist in the use of mixed data sources for non-AI analytical methods, the application of AI perceivably raises new considerations of increased speed of analysis, and the way in which AI models would generate predictions changing the way KYC checks would be processed.

2.3.2 The ongoing relevance of existing regulatory guidance and frameworks

This use case highlighted the various ways in which existing requirements can be relevant. First, both the MAS and PDPC guidance were relevant, along with requirements around data quality, the necessity of utilizing third-party data, and vendor risks assessments that would mitigate inherent risks. Data protection regulation is important in the context of AI-facilitated KYC, especially given multiple sources of data have been leveraged, some of which are sensitive in nature. FIs would need to consider whether consent provided for the data to be collected and processed (especially those sourced from third parties) would allow it to be used in their AI models and systems. There could potentially be a basis for lawful processing of personal data based on 'legitimate interests'

2.3.2 The ongoing relevance of existing regulatory guidance and frameworks

In addition, General Data Protection Regulation (GDPR) requirements relating to automated decision-making processes were considered relevant. A data subject has a right to be informed when they are subjected to automated decision-making processes, and a right not to be subjected to a decision based solely on an automated process. This means that the data controller must provide data subjects with access to at least a right of human intervention, including potentially contesting a decision. Finally, any regulations that restrict cross-border data flows of personal data – for example, data localization requirements – would introduce complexities for operation of the AI system for KYC.

2.3.3 Fairness and how biases can be exasperated and mitigated

Ensuring the fairness principle was implemented effectively in this use case was a key consideration. Because the KYC checks were undertaken using mixed data sources as an input for AI systems, there was a risk that bias could occur if those datasets lacked sufficient data, diversity, or are skewed towards certain characteristics.

Another consideration was whether any of the existing regulatory guidance remained relevant in addressing these identified risks of bias. **The PDPC Framework's guidance** requires due diligence in the selection of datasets for model development. If a more representative dataset were used, these relationships may not feature, and the AI system could generate more accurate predictions without increasing **FEAT-related risks**. Another relevant feature of the **MAS and PDPC guidance** was consistent monitoring for bias in AI systems and data was particularly important.

2.3.4 Potential trade-offs between transparency and the security of AI applications

A key finding in the context of this use case was the need to balance transparency and confidentiality. It was observed that increased transparency around the workings of an AI model and related data could result in greater security/ privacy risks. Excessive disclosure of how a model functions could potentially counteract the effectiveness of an AI application for KYC and leave it open to manipulation.

It is important to consider if any regulatory guidance would be helpful when balancing these transparency and confidentiality expectations. The PDPC AI Framework has recognized there could be scenarios in which it might not be practical or reasonable to provide information in relation to an AI algorithm, especially in the context of anti-money laundering (AML) or fraud prevention. As such, a more suitable outcome would be a disclosure to a potential customer that an AI system was being used in a KYC process, while managing expectations about the degree of information that could be disclosed.

3. LESSONS LEARNED FROM COLLABORATIVE INITIATIVES ON THE APPLICATION OF RESPONSIBLE AI PRINCIPLES

3.1 Materiality-based Approaches Focus Efforts on More Sensitive Use Cases

3.2 Regulations and Existing Risk Management and Governance Frameworks Remain Relevant

3.3 Collaborative Efforts Between the Developer and Deployer of AI Can Promote Transparency and Accountability

3.4 Diversity of Perspectives Ensures Responsible Use Principles Are Appropriately Applied

3.5 The Importance of Context Underscores the Need to Share Use Cases



LECTURE 03

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3. LESSONS LEARNED FROM COLLABORATIVE INITIATIVES ON THE APPLICATION OF RESPONSIBLE AI PRINCIPLES

Cross-industry collaborations to apply responsible AI guidance to real-world use cases can generate valuable learnings and insights, beyond what would be generated through work as individual institutions. When these learnings are shared, there is a greater understanding of how AI is being used in the financial services industry (and in other industries), allowing for the identification of good practices to ensure it is being used responsibly.

From Microsoft's own responsible AI journey, and through participation in the above initiatives, an overarching learning was the importance of context in evaluating the regulatory implications of an AI application. Not all responsible AI principles will be equally relevant in each context, and regulatory issues are best assessed based on the specific use of AI being deployed, and the jurisdiction and industry in which the AI application is being used.

3.1 Materiality-based Approaches Focus Efforts on More Sensitive Use Cases

One key learning is the value of adopting a materiality-based approach to assess the regulatory considerations when using AI. A materiality-based approach helps focus efforts and resources on more sensitive use cases, compared with approaches that treat all uses of AI the same. This will avoid undertaking detailed considerations of AI guidelines for every use case regardless of impact.

A materiality-based approach has been endorsed by **MAS and HKMA** in their respective responsible AI guidelines, where firms are encouraged to calibrate actions and requirements under their internal governance framework based on the materiality of the **AIDA applications**. AI systems of 'high risk' (e.g., remote biometric identification systems) will be subject to greater scrutiny, carrying a set of five obligations; whereas 'unacceptable' AI systems are proposed to be banned as being contrary to **EU fundamental rights**.

Ch. 4- Responsible AI: An opportunity to test regulatory principles**3.1 Materiality-based Approaches Focus Efforts on More Sensitive Use Cases**

There are various resources that can be used for guidance in making this determination, including MAS guidance and the PDPC AI Framework's probability–severity of harm matrix. Some of the relevant considerations drawn from these resources include:

- The extent to which AIDA is used in decision-making;
- The complexity of AIDA model;
- The extent of automation of process of AIDA-driven decision-making;
- The severity and probability of impact on different stakeholders, including individuals;
- Monetary and financial impact;
- Regulatory impact; and
- Any options for recourse available.

3.1 Materiality-based Approaches Focus Efforts on More Sensitive Use Cases

In practical terms, FIs could apply a baseline set of responsible AI criteria to any AI system being built, with a more intensive review procedure in place for sensitive uses where there is greater risk of harm. Additionally, organizations may find benefit in referring to published use cases that apply responsible AI guidance, to assess how that organization has labelled an application in terms of materiality to inform where their own AI application would fall on that spectrum.

The second step is that once a determination of materiality has been made for a particular use case, an organization must consider the proportionality of the responsible AI use principles and measures against that determination. Different approaches are possible, including applying a subset of principles to a use case based on an assessment of what was relevant; or applying each principle but doing so in a proportionate manner for the specific use case. In some cases, there may be a finding that a particular principle may not be relevant to a use case, and organizations should be granted the discretion to make that determination.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.2 Regulations and Existing Risk Management and Governance Frameworks Remain Relevant

Another learning is that existing regulatory frameworks can be leveraged to assess regulatory principles and ethical implications of using AI in specific use cases. FIs already operate in a highly regulated industry – with most jurisdictions having robust regulatory requirements in place covering outsourced technology use, risk management, and governance. These existing regulatory frameworks are therefore already guiding an FI's approach in adopting AI and are relevant in translating responsible use principles into practice. **Table 4.3 (p. 104) provides** a hypothetical example of some of the relevant questions involved in an AI context and noting relevant existing regulatory requirements in Singapore that are applicable.

An associated point is that established internal risk assessment and governance frameworks constitute solid starting points for FIs to implement AI in an accountable way. A global survey conducted by the Institute of International Finance found that most respondents are applying their existing model risk management or enterprise risk management (ERM) frameworks to AI applications, and several others are developing enhancements to account for new risks arising from AI techniques, such as control processes to mitigate against bias and discrimination of models.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.2 Regulations and Existing Risk Management and Governance Frameworks Remain Relevant

Applying existing governance processes would continue to make humans accountable for the decisions that AI systems make. While internal governance processes are relevant, they still need to be applied with the nature of the AI system, and its associated materiality, in mind. For example, the challenges posed in ensuring transparency and intelligibility of AI models (across all use cases), or through the wider use of mixed datasets as inputs for AI systems (as in Use Case 3) warranted careful consideration through existing risk management frameworks. For example, the MAS Veritas Consortium and the European Commission's Ethics Guidelines for Trustworthy AI emphasized the need to understand how existing internal risk management and governance frameworks can be leveraged to ensure responsible use of AI.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.3 Collaborative Efforts Between the Developer and Deployer of AI Can Promote Transparency and Accountability

A third learning from Microsoft's journey is the importance of FIs and technology partners working together to promote transparency and accountability. Both the technology partners that develop AI solutions, and the FIs that deploy the AI technologies procured from these third parties, play an important role in the responsible use of AI. For example, key concerns for AI developers will be the model-building process and ensuring AI solutions are robust and explainable. AI deployers that procure these technologies will alternatively be concerned about getting sufficient information to understand the model so that they are able to explain the model to customers, make informed choices in their use of the system, and be accountable to their customers.

A growing area of interest is in the complementary roles that these two stakeholders can play in promoting the responsible use of AI. AI deployers should be empowered to leverage on their technology partners' help in providing sufficient information to their customers. Technology partners may consider sharing technical and non-technical information about the capabilities and limitations of AI software being developed and filling any gaps between marketing and technical documentation where relevant.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.3 Collaborative Efforts Between the Developer and Deployer of AI Can Promote Transparency and Accountability

Notably, we have observed regulators encouraging collaboration among technology partners and FIs deploying AI technologies. Examples include the MAS Veritas Consortium that brought together industry participants to create a framework for FIs to promote the responsible adoption of AI and data analytics. Similarly, the European Commission's Ethics Guidelines for Trustworthy AI were prepared by members of the 'High-Level Expert Group on AI', which brought together prominent companies across sectors **like AXA, Microsoft and Bosch, as well as universities and government advisers.**

Another important area of collaboration can be seen in the effort by technology partners to increase the set of tools available to help ensure the responsible use of AI by organizations deploying the applications. For example, cross-industry partnerships and research efforts are contributing towards the growing availability of technical approaches relevant to the responsible use of AI. For example, the Partnership on AI is working on a project (ABOUT ML) to improve transparency through the creation of standardized documentation on machine learning models and datasets. These efforts indicate the importance of the complementary roles of both the technology partner and the financial institution in increasing transparency around AI applications.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.4 Diversity of Perspectives Ensures Responsible Use Principles Are Appropriately Applied

A fourth lesson emerging from Microsoft's own journey, as well as the initiatives we have been involved in, is that there is great value in bringing a diverse set of perspectives into discussions on responsible AI. A diverse set of perspectives helps to ensure that the responsible use principles are appropriately applied, and that potential concerns are identified and dealt with early.

At the highest level, diversity can come from cross-border discussions with regulators and governments in different jurisdictions, to share use case examples and approaches to implementing the responsible use of AI. **For example, Singapore's Personal Data Protection Act 2012 (No. 26 of 2012) and Hong Kong's Personal Data (Privacy) Ordinance (Cap.486)** are very similar in that they regulate the collection and use of personal data. Data protection risks may arise from the input of personal data into AI technology when potentially large amounts of personal data are collected for big data analytics purposes, and datasets comprising personal data are used in AI applications. Hence, the Singapore and Hong Kong regulators may wish to collaborate to share their local perspectives and experiences on how the respective data protection laws in their jurisdiction would address the data protection risks that may arise from the deployment of AI solutions.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.4 Diversity of Perspectives Ensures Responsible Use Principles Are Appropriately Applied

As an example, in the collaboration between Microsoft, Standard Chartered, Deutsche Bank, Visa and Linklaters described above, the diversity of the project group analyzing the three use cases ultimately proved to be very valuable. The input of the data scientists was particularly insightful when the AI solutions for each use case were being explained and what their capabilities (and limitations) were. Equally, regulatory, compliance, and policy perspectives were needed given the societal interest in responsible AI. At a firm level, diversity can also come from engagement with different teams within an organization itself, including the teams responsible for developing the AI solutions, as well as the legal, risk and compliance teams. While these cross-practice conversations may not have been necessary or natural just a few years ago, the principles-centric approach to AI adoption that requires a deep understanding of the model and its impact necessitates this type of collaboration.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.4 Diversity of Perspectives Ensures Responsible Use Principles Are Appropriately Applied

Microsoft takes this multidisciplinary approach to responsible AI governance with the **AETHER (AI and Ethics in Engineering and Research) Committee**, which leverages scientific and engineering resources to provide subject-matter expertise on the state-of-the-art and emerging trends regarding the enactment of Microsoft's responsible AI principles; the Office of Responsible AI, which sets policies and governance processes; and the **Responsible AI Strategy in Engineering (RAISE) group**, which enables engineering groups to implement responsible AI processes through systems and tools. It is important for a diverse set of views from people of different backgrounds to contribute to any decisions around the implementation of AI to ultimately ensure such decisions are reflective of the communities that organizations serve and operate within.

Ch. 4- Responsible AI: An opportunity to test regulatory principles**3.5 The Importance of Context Underscores the Need to Share Use Cases**

A final lesson is that use cases and examples can strengthen understanding across the industry of what regulatory issues may apply when using AI. As highlighted previously in this chapter, the context in which specific AI solutions are deployed is of utmost importance in evaluating the regulatory implications – and not all responsible AI principles are equally relevant in every context. There is a wide spectrum of AI use cases differentiated by not only functionality and purpose, but also the type of AI model used, model training (including the data that is used for that training), and impact on data subjects. There has been a shift in Asia towards practical considerations in translating principles to practice based on specific use cases, as the actual deployment of AI by financial institutions increases. In the past, most discussions around responsible use of AI in finance were centred around defining ethical principles and raising awareness of the need for responsible

AI.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

3.5 The Importance of Context Underscores the Need to Share Use Cases

Through this project, the partners worked collectively to assess the responsible AI considerations involved in three different financial services use cases: predicting travel patterns based on previous spending; automating verification of wet ink signatures; and regulatory compliance related to know your customer checks. By basing its work on specific use cases prepared by the financial institutions involved, together with technology company participants like Microsoft, the group is generating important learnings that strengthen collective efforts to implement AI responsibly. HKMA followed a use case-based approach reflected in its white paper on AI, where it shared use cases and potential applications to illustrate how banks can responsibly integrate AI into their business models. As more institutions share use cases where they show how they worked through the responsible AI considerations involved, the whole industry will benefit. There are similar efforts outside the financial services industry to gather and publish use cases for responsible AI in Asia. This includes the pilot of Australia's AI ethics principles (where Microsoft was invited as one of five companies to participate), and Singapore's two volumes of use cases compiled by the IMDA, a number of which relate to financial services.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

LECTURE 03

4. Conclusion

This chapter explore that despite the growing adoption of AI technologies in the financial services industry, and the increased focus on how these technologies can be responsibly implemented, to date regulators have proceeded cautiously in this space. Regulators generally have not issued new regulations that are specific to AI but have first focused on issuing non-binding guidance and principles. This principles-first approach has created an opportunity for collaboration between regulators and industry – with that dialogue allowing regulators to refine and converge on basic principles of responsible AI, as well as gain insights on whether new and specific regulation is needed to govern the use of AI.

Ch. 4- Responsible AI: An opportunity to test regulatory principles

LECTURE 03

Additionally, not all responsible AI principles will be equally relevant in each context, and regulatory issues are best assessed based on the specific use case, the AI application deployed, and the jurisdiction and industry in which the AI application is being used.

A final learning is that sharing use cases and examples can strengthen understanding across the industry of what regulatory principles and ethical issues may apply when using AI. As with any journey, it is important to engage in dialogue and share learnings along the way, as ensuring the responsible use of AI will be an ongoing process and not one where any institution can claim to have 'arrived' at the end.

Artificial Intelligence for Financial Technology

References:

Required Reference:

Remolina, N. and Gurrea-Martinez, A. eds., 2023. Artificial intelligence in finance: Challenges, opportunities and regulatory developments. Edward Elgar.

Recommended Readings

Chishti, S., 2020. The AI book: the artificial intelligence handbook for investors, entrepreneurs and fintech visionaries. John Wiley & Sons.

The background of the slide is a grayscale photograph of a person's hands typing on a laptop. Overlaid on this image are several semi-transparent icons: a graduation cap, a target with an arrow, an open book, a magnifying glass, a bar chart with a person presenting, and a group of three people. The title 'The End of Lecture' is written in a white, italicized serif font across the center of these icons.

The End of Lecture

Thank You