



afme/

asifma

sifma

26 October 2020

Mr. Alp Eroglu
Senior Policy Advisor
IOSCO General Secretariat
Calle Oquendo 12
28006 Madrid

via email: consultation-02-2020@iosco.org

Dear Mr Eroglu,

GFMA Consultation Response - IOSCO Consultation Report - The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers

The Global Financial Markets Association (GFMA) welcomes the opportunity to comment on the IOSCO Consultation Report (referred to hereafter as “the CP”) on **THE USE OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING BY MARKET INTERMEDIARIES AND ASSET MANAGERS**. The GFMA represents the common interests of the world’s leading financial and capital market participants, to provide a collective voice on matters that support global capital markets. We advocate on policies to address risks that have no borders, regional market developments that impact global capital markets, and policies that promote efficient cross-border capital flows, benefiting broader global economic growth.

The Association for Financial Markets in Europe (AFME) in London, Brussels and Frankfurt, the Asia Securities Industry & Financial Markets Association (ASIFMA) in Hong Kong and the Securities Industry and Financial Markets Association (SIFMA) in New York and Washington are, respectively, the European, Asian and North American members of GFMA.

Our high-level response to the CP is provided below and followed by a response to each question.

We and our members stand ready to engage on this topic further with IOSCO (and members of IOSCO). We look forward to having the opportunity to provide further assistance as Artificial Intelligence policies continue to be developed.

Respectfully,

Allison Parent
Executive Director
Global Financial Markets Association (GFMA)
aparent@gfma.org
www.gfma.org

Executive Summary:

The GFMA welcomes the opportunity to input into IOSCO's work on 'Artificial Intelligence' (AI) and 'Machine Learning' (ML). The use of AI/ML by market intermediaries and asset managers is growing but is, in many cases, still in very early stages. As the CP notes, there are significant efficiencies and benefits to be gained for both firms and investors. We therefore encourage IOSCO to support the global development of AI/ML within capital markets and avoid fragmentation and overregulation, which could slow down their adoption and development.

As a general principle, GFMA supports a technology-neutral regulatory approach, which is focused on outcomes rather than specific techniques or processes. This approach provides the flexibility needed for firms to implement the appropriate controls for the activities they are conducting in a risk-based and proportionate manner.

Furthermore, capital markets is already a highly regulated sector (a critical factor which we feel is not adequately discussed in the CP), with many of the existing requirements equally applicable to firms' use of AI/ML. For example, many jurisdictions impose obligations relating to areas such as conduct risk, duty to clients, internal governance, outsourcing/third-party risk management, technology, data privacy and model risk management. In light of this, we encourage IOSCO to survey for areas where gaps may exist, rather than broadly applicable requirements, and to consider focusing on AI/ML applications that may present a high risk to clients.

However, we appreciate that a globally harmonised regulatory approach to a complex and developing area such as the use of AI/ML in capital markets is also important, and will bring significant benefits for GFMA's members, many of whom operate cross-border. Fragmentation of approaches between jurisdictions adds additional cost and complexity for firms, which limits the potential benefits for both firms and their clients.

In relation to the specific questions set out in the CP, we have provided detailed answers below, but wish to highlight the following key points:

- **Definition of AI/ML:** a clear definition of AI/ML is key to ensure that the scope of the guidance is interpreted in a harmonised fashion. The definition should be accurate, future-proof and specifically exclude non-AI/ML techniques. We have provided a suggestion in our response below;
- **Proportionality of application:** firms should be able to apply a risk-based approach to the application of the guidance, taking into account a range of factors (such as the process to which AI/ML is being applied, the level of existing regulatory coverage and the impact on clients), not just the size of the firm;
- **Granularity:** The guidance should avoid being overly specific, as this risks a 'one-size-fits-all' approach, preventing firms from taking into account their existing structures and governance arrangements, as well as the specific controls of skills that will be needed in different situations;
- **Transparency:** as a key focus of this CP relates to transparency and disclosure, we note that level of transparency required for any AI/ML application will be dependent on the needs of the stakeholders involved and will therefore be highly variable. The guidance should allow firms to make a risk assessment and determine the appropriate level, rather than mandating a single standard.

Consultation Response:

1. Do you agree with the proposed definition of AI and ML?

'Artificial Intelligence' (AI) and 'Machine Learning' (ML) are often broadly defined in reports and academic papers which makes it difficult to apply these definitions in practice. For example, the 2017 Financial Stability Board (FSB) report 'Artificial Intelligence and Machine Learning in Financial Services', defines AI broadly as "the theory and development of computer systems able to perform tasks that traditionally have required human intelligence".¹ The World Economic Forum interprets the term slightly more narrowly, noting that AI "...develops computers that can do things traditionally done by people...sense or perceive the world and collect data... and act independently – all underpinned by the ability to learn and adapt over time".²

If used to develop policies, these definitions may create uncertainty on the scope of the technologies subject to the proposed regulatory framework.

Therefore, the challenge is how to provide practical definitions that are accurate (e.g. avoiding the inclusion of other non-AI analytics technologies), future-proof (e.g. considering the pace of innovation in the field) and broadly harmonised across major jurisdictions.

In relation to IOSCO's definition of AI, we suggest as a starting point the following drafting provided by AFME to the European Commission's on its 2020 White Paper 'Artificial Intelligence – A European Approach to Excellence and Trust':

*"Artificial intelligence (AI) systems are systems that act in the physical or digital world by perceiving their environment through data acquisition, interpreting the collected data, reasoning on the knowledge, or processing the information derived from this data and identifying the best action(s) to take to achieve the given goal. AI systems may adapt themselves or their own algorithms by analysing how the environment is affected by previous actions, knowledge or data."*³

This drafting, which built on an earlier suggested definition⁴ from the European Commission's High Level Expert Group (HLEG)⁵ on AI, was based on the reasoning that a definition of AI for policy purposes:

- should be concise and avoid unnecessary explanation of well understood terms;
- should not include anything that is not an essential and defining characteristic of AI;
- should focus on the ability of AI systems to perceive, reason, interpret, decide, learn and adapt;
- should not list types of AI models, in case this is unintentionally interpreted as an exhaustive list; and
- should not suggest that AI systems necessarily make or action the ultimate decision - AI identifies the most efficient solution to achieve a goal, but may or may not have the ability to execute.

Clarification that the scope of this guidance specifically excludes Robotic Process Automation (RPA) and any rules-based algorithms (i.e. where AI/ML is not utilised) would also be helpful.

In relation to the definition of ML, we agree that it should be seen as a subset of AI (noting that the way that the definition of AI in the CP is worded appears to confuse this distinction, suggesting that AI always necessarily includes ML).

Finally, we would like to raise a broader point of scope, regarding the lack of a level playing field between regulated firms that will be subject to this guidance (and the broader regulatory environment) and firms which, although acting within capital markets, fall outside the scope of regulated activities. Not only does this distinction place different standards upon firms conducting the same activities, it also limits the ability of competent authorities to satisfy themselves that the appropriate governance measures are in place for AI/ML across capital markets.

¹ <https://www.fsb.org/wp-content/uploads/P011117.pdf>

² <http://reports.weforum.org/digital-transformation/artificial-intelligence-improving-man-with-machine/>

³ https://www.afme.eu/Portals/0/DispatchFeaturedImages/20200612%20AFME%20EC%20AI%20CP%20Response%20-%20Final_.pdf

⁴ <https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines>

⁵ <https://ec.europa.eu/digital-single-market/en/high-level-expert-group-artificial-intelligence>

2. Do you see any risks or challenges around AI and ML which are not mentioned in the report?

Overall, the analysis of risks and challenges presented in the CP is relatively balanced, however we would like to raise the following points.

In relation to governance and oversight:

- It should be taken into account that, while senior accountability is important, oversight of technology such as AI is typically organised across the three lines of defence and that the responsibilities and necessary skills will differ across these roles. We return to this point in our comments regarding Measure 3 under Question 3 below.

In relation to the ethical themes (see also, AFME's White Paper 'Considerations on the Ethical Use of Artificial Intelligence in Capital Markets'⁶):

- Non-maleficence: While we agree that the aim of AI should be to 'do no harm', the definition of harm should be carefully considered, and that 'prevention of harm' may be a better principle to use. For instance, if a firm uses an AI application to perform suitability checks on its clients, it should not be considered harmful to withhold services from clients that do not pass the assessment. Indeed, the ramifications of providing unsuitable services to individuals or groups can be significant and harmful to society more broadly.
- Human autonomy: This is a broad term, which can be taken to mean different things. For example, the European Commission's HLEG on AI's use of the term in its Ethics Guidelines for Trustworthy AI encompasses the rights of subjects of an AI application to have "full and effective self-determination over themselves".⁷ By distinction, the explanation given "*humans have power over what the model can and cannot decide*" would be more accurately called human 'control' or 'oversight'. We suggest that one of these terms is used instead to avoid confusion. Furthermore, there are situations in which (while human oversight of the application remains overall, including with suitable guardrails and other controls) it may be extremely important for a human interacting with an AI application to be subordinate to that application. For example, it may be necessary to prevent the over-riding of AI applications concerned with safety systems or the detection of crime, such as anti-money laundering (AML) applications. In such a situation, the decisions made about the levels of control and human interaction with the application would be suitably documented with an auditable maker/checker process. We suggest that this theme is therefore amended accordingly to "*human control*" or "*human oversight*".
- Explainability: AFME's White Paper 'Artificial Intelligence and Machine Learning in Capital Markets: Considerations for a Broad Framework for Transparency'⁸ focused on transparency as helpful in meeting stakeholder needs, rather than the narrower term of explainability, which we understood as the extent to which the complex internal mechanics or workings of a model can be expressed.⁹ The paper also challenged the often restrictive or binary way in which explainability is discussed: that a model is either explainable or it is not. This expression of explainability does not allow for the spectrum of AI/ML models that can be used, each of which will have different features and levels of explainability. Equally, it does not allow for developments in explainability techniques and understanding.

We are therefore concerned by the way in which explainability is expressed and the fact that it is included under the ethical themes of the CP: "*Explainability: ensuring the outcomes arising out of the models used can be explained*". This is too complex a subject to be expressed so concisely – it risks (i) conflating technical explainability with broader transparency and (ii) suggesting that technical explainability is always required.

Instead, we suggest that transparency as outlined in our White Paper may be a more helpful concept and we support the statement on page 12 that "the level of transparency will also differ depending on the audience; for

⁶ [https://www.afme.eu/Portals/0/DispatchFeaturedImages/afme-cao-ai-considerations-on-the-ethical-use-of-ai-30-nov-2018%20\(1\).pdf](https://www.afme.eu/Portals/0/DispatchFeaturedImages/afme-cao-ai-considerations-on-the-ethical-use-of-ai-30-nov-2018%20(1).pdf)

⁷ <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>

⁸ <https://www.afme.eu/reports/publications/details//AITransparency>

⁹ On explainability, see H. Hagras, "Toward Human-Understandable, Explainable AI", *Computer*, September 2018

example, a regulator may require more detailed information than a client. These considerations need to be balanced in determining the appropriate level of transparency in the use of AI and ML.”

We would also like to raise a strong concern that Chapter 5, which considers firms’ responses to the potential risks arising from the use of AI and ML, did not discuss the mitigants such as technological and financial processes and procedures. In doing so, the CP assumes, as a base case, that there are no existing regulatory requirements in place to mitigate the risk of AI models. However, as capital markets is a highly regulated industry, firms (particularly those subject to prudential regulation) already use a wide range of controls, from the Model Risk Management required by many regulators, to broader lifecycle management of AI/ML applications. We encourage IOSCO to consider the regulatory mitigants that are already in place to address the risks that are relevant to AI and, if gaps are identified (for example in the use of AI by firms that are not regulated entities within financial services), that this is addressed in a targeted fashion, so as not to conflict or duplicate existing regulations.

3. Do you agree that the guidance set out in Chapter 6 of the Consultation Report is appropriate to address the potential risks associated with the general use of AI and ML by market intermediaries and asset managers? If not, please provide details.

The GFMA would like to raise the following in relation to the guidance:

Proportionality

- We are strongly supportive of the principle of proportionality in the application of the guidance and request that it is made one of the Measures in its own right. We agree that size should not be considered as a relevant factor when considering proportionality. Just as the CP suggests in the example provided for small firms, the use of AI/ML by big firms does not necessarily have a material impact on market participants and should not by default be subject to new requirements. The impact on client outcomes and markets - not the size of the firm - should be the key factor considered when applying this principle. In this sense, the use of AI/ML, e.g. for market making purposes (placing quotes in the markets), should not be subject to the same requirements as trading algorithms, for example. It could even be the case that for certain applications it is not necessary to establish new requirements when they do not pose a risk to the client. For example, this could be the case when using AI/ML for marketing purposes.
- The GFMA believes that regulators and policymakers should take a technology-neutral approach to rulemaking and the supervision of financial markets. Therefore, we believe that designing regulatory frameworks for specific technologies is not an ideal solution and that the focus should instead be on regulating outcomes rather than the technology used. This allows the rules to be designed in a principles-based manner that applies to activities rather than the means used.
- Capital markets is a highly regulated industry. Where regulatory frameworks do exist for particular technologies, risk-based application is critical. The materiality of the use case should be assessed, as well as whether the use of the technology brings new, or changes existing, risks. Where the risk profile remains the same as when the technology is not used, the regulatory framework should not seek to impose additional requirements, which could otherwise overlap with existing requirements and create inconsistencies.

Measure 1 (focused on senior management oversight):

- Strong governance and oversight of the use of technology is a key principle for GFMA’s member firms. In many jurisdictions specific requirements exist to ensure that management have full coverage of the firm’s activities as well as the appropriate skillsets to perform their oversight.¹⁰ We therefore suggest that it is clarified that the requirement to have “*designated senior management responsible for the oversight of the development, testing, deployment, monitoring and controls of AI and ML*” may be covered by senior management already designated for the oversight of technology more generally. In addition, existing senior individuals/groups (e.g. model risk

¹⁰ As well as the UK regime noted in the CP, examples include the Hong Kong ‘Manager in Charge’ regime, or the Australian Banking Executive Accountability Regime (and proposed Financial Accountability Regime)

management) already exist to provide oversight of AI/ML models, so a designated AI/ML specialist may not be needed, provided that existing functions have the appropriate skillsets.

Measure 2 (focused on testing and monitoring):

- The notes for Measure 2 states that *“Once deployed, the performance and output of the AI and ML algorithms should be monitored on a real-time basis.”* However, real-time monitoring may not be necessary, depending on the use case, or even practical.
- Similarly, the notes state that firms should ensure that adequate *“kill switch” functionality is built into their control framework, and that appropriate functions and people have access to it* as well as backup solutions supporting this functionality. While kill switches may be appropriate for some AI/ML applications, such as when AI/ML is used for algorithmic trading, they may not always be appropriate for other use cases. We suggest that the wording is amended to *“appropriate controls”* which may include, for example, guardrails/caps/floors etc.
- More generally, as noted in our response to Question 2 above, Chapter 5 of the CP listed the responses to the potential risks arising from AI/ML under culture, accountability, knowledge/expertise/skills and operational resilience, and did not discuss mitigants such as technological and financial processes and procedures. As capital markets is a highly regulated industry, firms already use a wide range of controls, from the Model Risk Management required by many regulators, to broader lifecycle management of AI/ML applications.

Measure 3 (focused on skills and expertise):

- We would appreciate clarification that what should be considered *“appropriate skills, knowledge, experience and expertise”* will naturally differ between functions and roles. For instance, even in the examples listed *“front line, compliance, risk management and senior management”*, it will not be necessary for all staff working on a firm’s use of AI to have in-depth technical knowledge provided that they are able to draw on this from elsewhere within the organisation as needed.
- While governance and structure naturally differ between firms, we note the importance of maintaining a distinction between the Business, which should have primary responsibility for due diligence of any AI/ML applications that it deploys, and its Control Functions, which should maintain an oversight, challenge and/or advisory focus as appropriate.
- Finally, we note that it is also crucially important that regulatory bodies develop the skills and resources to respond to and support the development of AI within their industries. This will also allow development of AI as a regulatory tool, for example for assessing large quantities of data or predicting the build-up of risk.

Measure 4 (focused on third party providers):

- We suggest that *“sanctions”* for poor performance, in an outsourcing context, may be misunderstood given the broader meaning of that term. Instead, we suggest using the term *“rights and remedies”*.
- Overall, we suggest that this area would be better handled contractually between the parties, with the necessary auditability provisions in place to ensure verifiability and enforceability. We note that separate guidelines exist in many jurisdictions for the management of outsourcing and third-party relationships, and that IOSCO itself has recently consulted on its Principles on Outsourcing¹¹. The GFMA has also provided comments to this consultation, which focus on the strong need for international harmonisation in regulatory approaches; a clearly defined scope which leverages established definitions; further consideration of data access and data localisation; and dialogue between regulators and the industry regarding measuring and addressing concentration risk.¹²

Measure 5 (focused on AI/ML disclosure):

- The term *‘sufficient’* for information disclosure is fairly broad and therefore likely to be unhelpful. As discussed in AFME’s White Paper ‘Artificial Intelligence and Machine Learning in Capital Markets: Considerations for a Broad Framework for Transparency’, the level of transparency required for any AI application will be dependent on the needs of the stakeholders involved and will therefore be highly variable. Firms must use a risk-based approach to determine what should be disclosed, to whom and how, in order to meet these needs and avoid

¹¹ <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD654.pdf>

¹² https://www.gfma.org/wp-content/uploads/2020/09/gfma-response-iosco-principles_on_outsourcing_cp.pdf

unnecessary confusion or mistrust. In many cases, it may be sufficient to include a simple disclosure to the effect that a particular process uses AI/ML. For this reason we, strongly advise that Measure 5 should be reworded as follows “Regulators should permit firms to consider what level of disclosure of the use of AI and ML, is required depending on the use case and nature of the data involved ~~by firms~~”, in order that firms may use a risk-based approach to transparency and avoid a one-size-fits-all obligation.

- We note that there is no mention in the Measure of complying with existing regulatory obligations, for example the rights of individuals in relation to automated decision making provisions under the EU General Data Protection Regulation (GDPR), and suggest that this is taken into account in the text.

Measure 6 (focused on controls for bias):

- While we agree that datasets may contain biases, it should not be assumed that this is inevitable and/or that complete removal of all bias is a prerequisite for the use of such data within an AI application. Instead, such risks should be mitigated, taking into particular account any risks to clients, including via ongoing assessment of the application’s outputs. We therefore suggest that the drafting is amended to “*unjustifiable bias*” or “*unfairly biased outcomes*”.
- Furthermore, bias is not necessarily always undesirable. For example, financial institutions have the obligation before concluding a credit agreement, to make a thorough assessment of the client’s creditworthiness, taking appropriate account of factors relevant to verifying the prospect of the client to meet his/her obligations under the credit agreement. Similarly, the use of AI/ML in marketing of products to clients or suitability assessments should avoid unfairly biased outcomes but should not prevent firms from being able to distinguish between different clients’ situations and needs.
- The expectation should be that the AI/ML application is not increasing the disparate impact compared to the training data i.e. it is at least as fair as the data used. Expecting AI/ML to reduce bias compared with human decision making should not be a policy requirement as it’s not always achievable.
- The notes for Measure 6 state that “*firms should ensure that the data set is representative of the target population so that they do not lead to exclusion phenomena*”. We suggest that it is more feasible and appropriate for firms to make efforts to determine whether a data set is unrepresentative, rather than prove that it is representative. Taking this into account, we also note that it is general practice to tailor the AI/ML application usage and limitations in decision making to the data used.

4. Do you disclose information to customers / clients on the use of AI and ML? If yes, please indicate what kind of information is disclosed?

GFMA is a trade association, and as such cannot comment on the individual practices of its members. However, as noted in our comments under Measure 5 above, and our response to question 5 below, when firms use AI/ML, consideration is given to client disclosures. In some cases, this is driven by regulatory obligations such as the GDPR in the EU, but it will also be driven by a risk-based assessment of the level of transparency required.

We note the example given on page 12 of the CP of a “black box” AI model; however, as we have discussed under Question 2 above, the use of such binary language in relation to technical explainability is restrictive and does not allow for the variety of AI/ML models and explainability techniques that now exist or may be developed in future. Furthermore, it conflates technical explainability with the broader and often more helpful concept of transparency.

5. What factors do you need to take into account when considering the appropriate level of information that should be disclosed to clients (including prospective clients) and other relevant stakeholders around the firm’s use of AI and ML algorithms?

As discussed in AFME’s White Paper ‘Artificial Intelligence and Machine Learning in Capital Markets: Considerations for a Broad Framework for Transparency’, we suggest that it is important first to determine which stakeholders (which may include clients) will require a level of transparency in relation to the firm’s use of a particular AI application, and what needs they are driven by. For example, transparency may:

- contribute to the feasibility analysis and business justification for a AI/ML project;
- assist in the assignment and monitoring of accountability;
- provide a suitable basis for the firm's management to sign off and oversee the firm's use of AI/ML;
- enable the firm's technology function to develop, monitor and optimise AI/ML applications;
- give assurance to internal users of the application of its benefits and performance;
- allow the firm to address concerns that external users or data subjects may have about their interaction with the firm's AI/ML application(s);
- demonstrate compliance with ethical and regulatory obligations; and
- enable oversight, auditability, and challenge by control functions (e.g. compliance, risk and internal audit).

The level of transparency can then be tailored to the individual needs of each type of stakeholder (e.g. for some stakeholders, technical explainability may be required, whereas for others, information about the processes and procedures surrounding the development of the AI/ML application may be more appropriate). For those interacting with the application, for example a firm's clients, it may be sufficient in some cases to use more generic disclosures, for example to notify them that AI/ML may be used in data processing, without going into detail, particularly where there is no demand for further information. Alternatively, client disclosure may happen through notice or consent in client/product terms and conditions (T&Cs). This generally uses comprehensive, standard language, as making changes to T&Cs can be time-consuming and expensive.

The AFME paper suggested how transparency could be provided in relation to the assumptions used in the AI/ML model development and the testing to which it was (and will continue to be) subject.

Finally, given the complexity of this topic, we note that some firms may wish to engage in further dialogue with their local regulators regarding the approach they are implementing in relation to transparency and disclosures.

6. How do you consider and apply proportionality to your systems and controls over the use of AI and ML?

The GFMA agrees that proportionality should be a key principle for regulators and supervisory in their oversight of AI/ML use, and for firms in their development and deployment of AI/ML. Tailoring requirements to the risk profile of the activity and its potential impact is the most suitable way to ensure that requirements are applied in a proportionate manner.

As part of firms' risk assessments, the following factors may be considered:

- The materiality of the activity to which the technology is being applied;
- The complexity of the model;
- The expected benefit from the use of the technology;
- The risks which are already applicable to the activity prior to the use of the technology, and the extent to which the use of the technology will change, mitigate or increase any potential negative impacts of the technology;
- The regulatory obligations, industry standards and internal policies or procedures to which the activity or technology is subject;
- The types of stakeholder who will be involved or affected, and in particular any impact on clients; and
- The monitoring, testing and overall lifecycle management to which it will be subject.