CHARTERED PROFESSIONAL ACCOUNTANTS CANADA
COMPTABLES PROFESSIONNELS AGRÉÉS CANADA

Chartered Professional Accountants of Canada
277 Wellington Street West Toronto ON CANADA M5V 3H2
T. 416 977.3222 F. 416 977.8585
www.cpacanada.ca

Comptables professionnels agréés du Canada
277, rue Wellington Ouest Toronto (ON) CANADA M5V 3H2
T. 416 977.3222 Téléc. 416 977.8585
www.cpacanada.ca

October 26, 2020

Mr. Alp Eroglu
Senior Policy Advisor
International Organization of Securities Commissions (IOSCO)
Calle Oquendo 12
28006 Madrid
Spain

Dear Mr. Eroglu:

Public comment on the use of artificial intelligence and machine learning by market intermediaries and asset managers

Chartered Professional Accountants of Canada (CPA Canada) is pleased to respond to the International Organization of Securities Commissions' (IOSCO) consultation report *The use of artificial intelligence and machine learning by market intermediaries and asset managers* (Consultation Report).

Data and artificial intelligence (AI) governance are key areas of focus for the Canadian accounting profession. Realizing the potential and benefits of AI and machine learning (ML) will require creating a framework to ensure human-computer interactions help and protect the public interest. CPAs are in a unique position to help define the principles and controls to guide AI design, development and deployment in a way that demonstrates integrity, transparency and accountability. CPA Canada is actively involved with the CIO Strategy Council[1] and Data Governance Standardization Collaborative[2], which develop standards to support emerging technology. CPA Canada has also published guidance on AI/ML[3], data governance[4] and cybersecurity[5] which support many of the measures outlined in the Consultation Report.

We are supportive of efforts by IOSCO to address the pervasive risks associated with expanded global use of big data and AI by investment firms. As noted in the Consultation Report, several national and regional securities market regulators have already issued or are in the process of developing regulation and guidance for their markets. It is important for the efficient operation of global capital markets that such AI and ML regulation/guidance is consistent.

In formulating our response to the Consultation Report, we consulted with a wide range of stakeholders including business leaders, investors, auditors, regulators, academics and experts from Canada's fintech and AI research communities, including the University of Montreal (Algora Lab), OBVIA, Fin-ML (Machine Learning in Finance) Network, MILA and Finance Montreal. Our detailed responses to select questions are included in the Appendix to this letter.

[1] https://ciostrategycouncil.com/standards/

[2] www.scc.ca/en/flagships/data-governance

[3] www.cpacanada.ca/en/business-and-accounting-resources/other-general-business-topics/information-management-and-technology/publications/tech-resources-for-cpas

[4] www.cpacanada.ca/en/foresight-initiative/data-governance/mastering-data

[5] www.cpacanada.ca/en/business-and-accounting-resources/other-general-business-topics/information-management-and-echnology/publications/questions-directors-should-ask-about-cybersecurity

For the proposed regulatory framework to have maximum effect, we believe the following key areas need to be addressed:

1. **Data and AI standards initiatives**

   The Consultation Report comes amid multiple guidance propositions and existing guidelines from other national and international bodies working to regulate the use of AI and ML, however, these are not referenced in the proposed regulatory framework. For example, the CIO Strategy Council has developed a standard *Ethical Design and Use of Automated Decision Systems*[6] that touches on all six measures listed in the Consultation Report.

   Any proposed regulatory framework on the use of AI and ML will require robust data standards that apply to companies regardless of their geography, size, and industry. There is an opportunity for IOSCO to take a more active role in relevant data/AI standards development initiatives and leverage existing frameworks and standards where possible.

2. **Clarification of key definitions**

   The AI space is evolving rapidly. The proposed definitions of AI and ML could be improved by broadening the definitions to include a focus on the future evolution and use-cases for AI. More specifically, structured data and the predictive capabilities of AI were well represented in the Consultation Report, however, the use of unstructured data to generate new data and insights by AI and ML was not equally considered. We believe the limited definition could create a blind spot when identifying potential risks of using AI and ML. Therefore, we suggest exploring other widely accepted definitions for AI and ML that are available or consider the alternative definitions we provide in the Appendix.

   In addition, we believe foundational high-level definitions of important concepts such as governance, oversight, bias and fairness should also be included. A common understanding of these terms will make it easier to understand and apply the proposed recommendations. We believe establishing a connection between the risks listed in Chapter 4 and the measures put forward in Chapter 6 would also be helpful.

3. **Ethics of AI**

   In addition to complying with all relevant laws, regulations and guidelines, users of AI and ML also have ethical and social responsibilities. Trust in AI systems does not depend exclusively on the moral conduct of individual agents, but on the ethical structure of collective agents (organizations, firms) and systems. The Consultation Report does not address the ethical challenges posed by the development of AI and ML systems.

   Ethical principles should be considered in the context of organizational rules, a governance framework and machine settings.

4. **Data integrity and role of assurance**

   Data integrity is an important element of AI governance that should be addressed in the regulatory framework. Key data integrity attributes to consider include relevance/fitness, completeness, accuracy, validity and currency of the data. These key attributes should also apply

[6] CAN/CIOSC 101:2019 “Ethical Design and Use of Automated Decision Systems”. Retrieved at: https://ciostrategycouncil.com/standards/implement-standards/

to associated metadata, which provides contextual information for the data. Frameworks for obtaining assurance about data integrity have been published by CPA Canada and the AICPA.[7]

The proposed regulatory framework should also recognize and consider how to apply assurance to the AI supply chain. There is increasing reliance on third-party AI systems, however, there is a lack of specific guidance addressing the issues of accountability, transparency, and explainability when it comes to outsourcing to these third-party AI systems. The AI supply chain and outsourcing risks are more complex than articulated in the Consultation Report.

The Consultation Report relies too heavily on service-level agreements (SLAs) as a means to manage and provide oversight of outsourced services. The degree of compliance by the service organization with its SLA may not be known without an audit of the service organization's controls over its systems. Right-to-audit clauses and requirements for assurance reports should be incorporated into SLAs. The AICPA has published guidelines and criteria for assessing service organization systems and controls.[8]

5. Cybersecurity and privacy risks

We believe there should be broader consideration for how the proposed regulatory framework intersects with privacy and cybersecurity risk management programs. The AICPA has published guidelines and criteria for assessing cybersecurity risk management programs.[9]

6. Mechanisms for customer and client feedback

While we agree that firms should disclose meaningful information to customers and clients around their use of AI and ML and the impact on client outcomes, it is equally important to have mechanisms for customers and clients to provide feedback. This would enable them to not only have a say in the level of information they require, but also allow them to dispute an AI-driven decision and to choose a human operator for decision-making.

7. Scalability of regulatory framework to smaller firms

While it is important to have consistent regulations across the industry, some proposed measures may be more challenging for smaller firms to follow. It is not clear in the Consultation Report whether special considerations would be given to smaller firms or if all firms would be expected to apply the same requirements.

A principles-based approach to the regulatory framework may be most effective in providing necessary guidance while preserving flexibility to tailor based on size or other risk-based factors.

8. Impact of AI technologies on sustainable finance

There is increasing recognition of the role that sustainability plays in a resilient financial system. Since the UN put forward Sustainable Development Goals in 2015, there has been a growing global awareness and urgency to shift investments to help meet these goals. This is changing the

[7] Criteria for Evaluating the Integrity of a Set of Data available at www.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/asec-data-integrity-criteria-ed.pdf

Describing a Set of Data and Evaluating its Integrity available at www.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/describing-a-set-of-data-and-evaluating-its-integrity.pdf

A Framework for Information Integrity Controls available at www.cpacanada.ca/en/business-and-accounting-resources/audit-and-assurance/internal-control/publications/framework-for-information-integrity-controls.

[8] Report on Controls at a Service Organization Relevant to Security, Availability, Processing Integrity, Confidentiality or Privacy available at www.aicpa.org/interestareas/frc/assuranceadvisoryservices/serviceorganization-smanagement.html

[9] SOC for Cybersecurity available at www.aicpa.org/interestareas/frc/assuranceadvisoryservices/aicpacybersecurityinitiative.html

way asset owners, managers and intermediaries assess, manage and report on environmental, social and governance (ESG) performance. It is important the Consultation Report recognize the transformational role AI technologies will be called on to play in this area in the coming years. In our opinion, this important theme is under-represented in the Consultation Report. We believe an in-depth consultation is necessary to consider the material association and influence of AI technologies on the broader sustainable finance discussion and the related regulatory interventions required.

9. **Applicability of the framework beyond the investment industry**

   While this consultation is specifically focused on asset managers and market intermediaries, we believe that similar considerations would apply to public companies. It raises questions about whether existing regulatory requirements around internal controls for public companies are adequate to manage the risks associated with the expanded use of big data and AI. We believe this warrants consideration by IOSCO.

We wish to acknowledge the generous contributions of a core group of AI researchers with whom we closely collaborated in preparing our response. The names of the individual members are included in the Appendix.

We appreciate the opportunity to participate in this important consultation. We highlight that the objective should always be to balance measures designed to mitigate potential risks while allowing investment firms to compete effectively and facilitate the best outcomes for investors.

We would welcome the opportunity to discuss these comments in greater detail and answer any questions you may have related to them. We would be pleased to convene a roundtable and provide additional support as you progress this important work.

Please contact Rosemary McGuire, Director, Research, Guidance and Support (rmcguire@cpacanada.ca) or Michael Wong, Principal, Research, Guidance and Support (michaelwong@cpacanada.ca) if you have any questions regarding our response.

Sincerely,

Charles-Antoine St-Jean

Charles-Antoine St. Jean
President & CEO
Chartered Professional Accountants of Canada

Appendix: Detailed Response to Consultation Questions 1, 2 and 3

In preparing our response, CPA Canada collaborated closely with a core group of AI researchers. We would like to acknowledge members of this core group and thank them for their contributions to this response:

- Marc-Antoine Dilhac – Professor at University of Montreal, Canadian Institute for Advanced Research (CIFAR) Chair in AI ethics at Mila (Quebec Artificial Intelligence Institute), co-responsible for Deliberation on AI at the International Observatory on the Societal Impacts of AI (OBVIA)
- Anne-Marie Hubert – Chair of the Human Technology Foundation (Montreal)
- Rheia Khalaf – Director of Collaborative Research and Partnerships at Fin-ML
- Manuel Morales – Professor at University of Montreal, Associate Researcher at International Observatory on the Societal Impacts of AI (OBVIA)

<u>Question 1:</u> **Do you agree with the proposed definition of AI and ML?**

We believe the proposed definitions lack a layer of complexity necessary to serve as the building blocks of a comprehensive guidance report that would assist IOSCO members in creating effective AI and ML regulatory frameworks for market intermediaries and asset managers.

In our opinion, a shorter, more concise definition of AI and ML would be more accurate, and also encompass and give equal weight to two broad categories of applications:

1. applications that create value or more efficient processes through pattern recognition and predictive algorithms that automate and/or optimize investment decision processes;
2. applications that leverage alternative unstructured information to assist with the various dimensions of an organization whose main activity is focused on managing financial assets.

We believe the proposed definitions do not fully explore the branches of AI that deal with structuring data, retrieval of information, knowledge-based structures and ontologies core to the second category of AI applications outlined above. Considering the objectives of the report, this could be problematic as it creates a bias that permeates the report and that bends towards risks mainly associated with the first category of applications. Even though the Consultation Report mentions the second category of applications, the overall language and spirit of the document seems to be the result of reflections mostly centered around the first category.

We believe this bias could be problematic to the extent that it creates a blind spot when seeking to identify potential risks brought about by the use of AI and ML particularly in the growing area of environmental, social and governance (ESG) investing. Asset owners and investment managers are organizing themselves to be able to measure and report on their ESG performance. Market intermediaries and asset managers are under pressure to leverage AI and ML technologies to capture the relevant information required for ESG investing purposes. These fall under the second category of applications of AI and ML that we believe are under-represented throughout the document.

Moreover, in terms of consistency and accuracy of the definitions, we do not agree with the statement "…(AI) can be understood as a combination of mass data, sufficient computing resources and ML". The first two elements are common across a lot of different technologies, which means ML is the only differentiating factor for AI and yet ML is then defined as a subset of AI, not a prerequisite. This can be rephrased in the context that AI has experienced a resurgence due to these factors, but not that they "define" AI.

We offer the following alternative definitions that address the concerns raised above:

Artificial intelligence is a field of science that studies and tries to reproduce the different mechanisms that constitute human intelligence. This includes neuroscience, psychology, behavioural sciences, biology, anthropology, mathematics, statistics, engineering and computer science. AI also encompasses applied branches of these subject areas that attempt to reproduce human cognition. Such efforts are a concerted combination of computer science and statistical methods that exploit massive data sets and exponentially growing computer power.

Machine Learning is a field of computer science that focuses on designing algorithms and methods that effectively compress knowledge into a computer system in such a way that it can perform complex tasks through a process akin to "learning" as opposed to hard static programming. These methods are based on substantial amounts of data from which the system obtains information relevant for the task at hand. Different types of "learning" have been developed for different tasks. Some methods can be described as replicating human reasoning and learning from experience.

Additional suggestions to improve the original definitions:

- The section on AI is also missing some interesting information, such as the cognitive functions it is trying to mimic: reasoning, knowledge representation, planning, learning, natural language processing, perception and the ability to move and manipulate objects.

- Sub-fields of AI other than ML are not introduced. The article "What is Aritificial Intelligence: Definition and Sub-fields of AI" provides a list of the cognitive functions but also the sub-fields. The section on AI should probably include some detail on both10.

- The section on ML is confusing. It first talks about analyzing and looking for patterns and then gets into inductive learning, without specifically focusing on the biggest difference between traditional analytic tools and ML. This is important because ML develops these patterns based on "its own experience" rather than being specifically programmed.

- It may also be helpful to state that AI is being developed on a spectrum of human involvement from complementing human decision-making to acting autonomously.

- The "adaptive" dimension of ML was not introduced. The report seems to imply that all AI systems continue to learn once in production but only if they are adaptive. This is not accurate. You can use ML learning techniques in development and testing and then put a static model into production.

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

We suggest a more consistent and comprehensive set of principles and criteria is needed to adequately address risks in financial markets. In addition, we believe the following risk categories deserve greater attention:

- market stability,
- ethical objectives,
- sustainable finance,
- data privacy.

[10] www.softwaretestinghelp.com/what-is-artificial-intelligence

1. Mitigating risks of market instability

Market stability is a necessary condition for the existence of financial markets and investment over time. In our view, the use of AI/ML systems must contribute to efficiency, robustness (including cybersecurity) and trust to support market stability.

1.1 Efficiency

Since AI/ML systems pose a wide range of risks, it is critical for market intermediaries and asset managers to justify their use to their clients in plain language. In particular, market intermediaries and asset managers must demonstrate that they achieve better outcomes for their clients by using AI/ML tools than without them. This condition is obviously counterfactual, but regulators (or self-regulating firms) can meet this requirement by implementing an accurate assessment protocol with control samples[11].

1.2 Robustness and resilience

Robustness[12] of AI/ML systems outweighs efficiency when it comes to ensuring market stability. If an AI/ML system lacks robustness, it is no longer important that it is theoretically efficient. The main risk associated with the deployment of AI/ML systems in financial markets is their potential unpredictability and their vulnerability to cyberattacks.

When processing massive amounts of data and writing their own mathematical formulas to solve problems humans can't, AI systems behave like "black boxes" we are not able to decipher. Even when we can anticipate the outcomes of an AI system under favourable conditions, it may be unpredictable under less favourable conditions and may not perform as expected.

Moreover, numerous examples show that AI/ML systems can be fooled by adversarial attacks altering and poisoning the data processed by algorithms. While AI/ML systems are developed to optimize investments, enhance risk management and detect fraud, adversarial attacks may turn those systems into useless or malicious tools rendering fraud undetectable and increasing systemic risks.

Finally, supercomputing machines are increasingly prone to computing failures, which raises concerns about safety and requires additional investment in monitoring and oversight. Different techniques are being developed to ensure better resilience of AI systems. Resilience must then be high on the list of criteria when asset managers and market intermediaries choose an AI system to manage investments.

Safety, security, and resilience issues should be addressed specifically in the Consultation Report and adequately tackled in the proposed measures.

1.3 Trust

Trust is key to social and market stability in general and to a beneficial deployment of AI/ML systems in the financial market. The risk of instability associated with a low level of trust in AI/ML ecosystems is high and must be taken seriously. Trust depends on four conditions:

- reliability of AI/ML systems,
- common ethical principles,
- alignment of ethical, social and financial objectives,
- oversight and governance.

[11] Measure 2 of the Document suggests such a protocol to check the counterfactual criteria: "The testing should be conducted in an environment that is segregated from the live environment prior to deployment".

[12] See for instance OECD, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449.

We have already discussed the conditions for AI/ML reliability, namely, efficiency and robustness. We will examine some of the remaining conditions below[13].

2. Mitigating ethical risks

Ethical principles that guide the development and use of AI/ML needs to apply to organizations, systems and devices. Trust in AI/ML systems does not depend primarily on the moral conduct of individual agents but on the ethical structure of collective agents (organizations, firms) and systems such as AI/ML applications and devices. AI ethics addresses both organizational and design issues.

The Consultation Report, in reference to the IOSCO's Fintech Network contribution, identifies five principles ("themes"):

- beneficence,
- non-malfeasance[14],
- human autonomy,
- justice,
- explainability.

Ethical principles must be implemented in the form of organizational rules, leading to a governance framework and ethical machine settings. The IOSCO framework may also benefit from a set of core principles similar to those included in the in the Montreal Declaration for the responsible development of AI (2018)[15] . These principles are reproduced below. It should also be noted that some aspects are already addressed to various degrees within the Consultation Report. The principles of privacy, solidarity, diversity, and environmental sustainability are notably absent in the Consultation Report

Principles	Application to financial market
Well-being	• Do not use AI/ML to harm investors. • Use AI/ML to promote the investors' best interest.
Autonomy	• Provide investors with the best information on the use of AI/ML to help them make good decisions and give informed consent. • Ensure that those responsible for the deployment and use of AI/ML systems within the firm have the appropriate level of skills and knowledge to understand the implications of their use (whether or not their design has been outsourced).
Privacy	• Protect personal data of the investors and ensure stringent confidentiality. • Ensure that personal data are used in accordance with the agreed purposes and prevent any misuse.
Solidarity	• AI/ML should help improve risk management and allow for a more equitable distribution of individual and collective risks among investors.
Democracy	• Make code for AI/ML systems and training data sets accessible to regulators and make them auditable by the relevant authorities. • Make code for AI/ML systems explainable to regulators and clients in accordance with different standards of explainability.

[13] For a holistic approach to trustworthy AI, see the report elaborated by the European Commission High-Level Expert Group on Artificial Intelligence, "Ethics Guidelines For Trustworthy AI", April 2019;

[14] Beneficence and non-malfeasance are two sides of the same token, namely the value of well-being.

[15] See Annex 1 of the Consultation Report, "Canada".

Principles	Application to financial market
Equity	• Ensure that AI/ML systems treat like cases similarly and do not favor one investor over another. • Ensure that the use of AI/ML does not give unfair advantage to certain market intermediaries and asset managers over others. • Prevent bias in both the algorithms and underlying datasets.
Diversity	• Ensure the environments supporting AI development are inclusive and reflect the diversity of societal norms. • Ensure diversification in the selection of investments to avoid potential underperformance of investments.
Prudence	• Test using simulations and small-scale rollouts before deploying AI/ML systems broadly. • Implement impact assessments and monitor continuously. • Disclose errors in AI/ML systems on a timely basis, especially when they impact the integrity of investors. • Promote AI-powered financial stewardship
Responsibility	• Always maintain a degree of human oversight when making decisions that affect the interests and integrity of investors.
Environmental Sustainability	• AI/ML infrastructure must not create unwarranted impact to the environment. • AI/ML systems in the financial sector should aim to reduce the environmental footprint of investments.

3. Including sustainable finance criteria in AI development

Reorienting capital flows toward sustainable investments has become a regulatory imperative in many jurisdictions and AI technologies are making ESG integration and socially responsible investing more accessible. For example, AI/ML systems can help (1) trace, assess and report the ESG reporting standards of investments according to the chosen metrics[16], and (2) integrate this assessment into automated or assisted decision-making. These types of AI/ML investment applications should be paired with algorithmic tools aimed at automating and optimizing decision-making processes when providing guidance to asset managers and market intermediaries.

In this context, it is important the Consultation Report recognizes the transformational role that AI technologies are already playing and will be called upon to play in this area in the coming years. Not considering the criteria of sustainable finance in the development and use of AI/ML in finance contributes to creating a series of risks that are overlooked in the Consultation Report. We encourage IOSCO conduct additional research into these types of investment applications and determine what changes may be necessary to the proposed framework.

4. Managing privacy of customer data

Almost all use of AI/ML require the collection and use of data. When using customer data, there are significant ethical considerations, and it is important to consider a customer's fundamental right to privacy. Therefore, we suggest guidance on the use of private data, as well as operational guidance to support

[16] Major financial institutions not only agree on the importance of ESG criteria to build sustainable finance, but also converge on common metrics. See World Economic Forum, Measuring Stakeholder Capitalism: Towards Common Metrics and Consistent Reporting of Sustainable Value Creation, September 2020.

organizations developing their own technology to ensure confidentiality is respected be included in the Consultation Report.

Question 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 proposed measures cover a wide range of concerns and provide a good starting point for regulating the use of AI/ML. We have highlighted in Question 2 additional areas that require consideration throughout the various measures.

We also believe it would be helpful to users of the framework if the risks identified in Chapter 4 were linked to the measures in Chapter 6.

Detailed comments on the six proposed measures can be found below:

Measure 1: Regulators should consider requiring firms to have designated senior management responsible for the oversight of the development, testing, deployment, monitoring and controls of AI and ML. This includes requiring firms to have a documented internal governance framework, with clear lines of accountability. Senior Management should designate an appropriately senior individual (or groups of individuals) with the relevant skill set and knowledge to sign off on initial deployment and substantial updates of the technology

Start at the problem statement identification stage:

- The problem statement identification stage should be included in the internal governance framework, even before any AI development is planned.
- Risks should be identified and assessed early. Controls should be put in place at the design and implementation phase and tested and monitored continuously throughout AI deployment and until end of life.

Design an inclusive governance framework:

- The designated senior management team responsible for the oversight of AI should include a mix of technical and non-technical individuals. Ongoing training and continued education should be provided to ensure employees are keeping pace with evolving technological risks.
- A diversified governance team can reduce the bias and the uniformity in the methods and approaches used, and can take into account the point of view of all the multidisciplinary stakeholders involved, both internal or external.
- Accountability is intricately linked to transparency and traceability of the dataset, processes and people that made the decisions during the AI lifecycle.

Accountability goes beyond legal requirements:

- Users of AI and ML also have ethical and social responsibilities. Even if some of these responsibilities are non-binding or are only prescriptive, not respecting them can lead to reputational risk. This is in addition to compliance with all relevant laws, regulations and guidelines.
- A clear and adequate “AI corporate culture” that encourages social and ethical responsibilities can be established by creating social responsibility programs, codes of conduct, and values/ethics statements for employees.

Design a governance framework that is a component of an Enterprise Risk Management (ERM) program:

- The governance framework should not be developed in a silo structure. Rather, it should be a component of an ERM program.
- Existing standards can be leveraged and enhanced to adapt to specific risks related to AI/ML. These risks should in turn be included in the aggregated risk assessment.

Consider alternative organizational structures:

- Different approaches to validate the governance structure can be considered. For example: third parties for independent review, in-house Centre of Excellence or an Ethical Standards Board.

Measure 2: Regulators should require firms to adequately test and monitor the algorithms to validate the results of an AI and ML technique on a continuous basis. The testing should be conducted in an environment that is segregated from the live environment prior to deployment to ensure that AI and ML:
(a) behave as expected in stressed and unstressed market conditions;
(b) operate in a way that complies with regulatory obligations.

Testing and monitoring should continue until end of life:

- The "testing environment" can prove to differ from the live environment. Therefore, it is important to stress that testing, as well as monitoring, need to be in place continuously throughout AI deployment and until end of life.
- The measure should be more reflective/inclusive of all types of AI, and in particular, data. The measure should extend through the lifecycle of data, from sourcing through to purging.

Consider risk tolerance levels:

- Although errors are undesirable outcomes of any model, some could be tolerated to a certain level before the "kill switch" functionality is activated. Testing and monitoring should be performed in relation to a predetermined level of error tolerance that the company is willing to accept. Different measures should be identified within the tolerance interval and beyond.
- The nature and significance of the harm associated with an incorrect outcome can vary, and should be estimated at the initial phases of the project.

Degree of human oversight:

- The level and depth of monitoring should go hand in hand with the degree of human oversight, as well as with the nature and significance of an error as mentioned above. In "human-in-the-loop" situations, (where human oversight is active and involved, and the AI is only providing recommendations or content), the monitoring exists by definition. More human oversight should be available in human-over-the-loop situations (where the human is in a monitoring or in a supervisory role, with the ability to take control when the AI model encounters unexpected or undesirable outcomes). Cases where there is no human oversight over the execution of decisions, and the AI system has full control without the option of human override should be avoided.

Identify a set of stress-testing scenarios:

- Point (a) refers to "stressed market conditions", which can cover only a subset of unexpected situations. In order to include a larger set of potential use cases, we recommend broadening this exercise to include sensitivity and stress-testing scenarios beyond the business-as-usual conditions (e.g., edge cases) to determine how the AI system behaves in unexpected situations and unfavourable environments.

Identify a response to monitoring:

- The response to monitoring of an AI system is not mentioned. A framework should be identified describing the different steps to be performed when an AI model is not behaving as expected. There should also be regular model tuning to reflect the adaptive nature of AI systems, changes in market behaviour or preferences, and changes in the collection and incorporation of new data into the training sets, which in turn can lead to model changes that should be deployed in the "live" environment.
- The measure seems to assume that a) monitoring leads to control and b) monitoring and change management are the same. Neither is the case. Changes to algorithms must be managed. Someone has to act on the results of the monitoring. Systems management and control are critical and should be recognized in the report.

Measure 3: Regulators should require firms to have the adequate skills, expertise and experience to develop, test, deploy, monitor and oversee the controls over the AI and ML that the firm utilises. Compliance and risk management functions should be able to understand and challenge the algorithms that are produced and conduct due diligence on any third-party provider, including on the level of knowledge, expertise and experience present.

Beware of overconfidence:

- Safeguards should be put in place to prevent overconfidence or overreliance on the AI/ML system. There should be risk oversight that determines the required level of task allocation between humans and automated processes.
- Mechanisms and measures should be put in place to ensure a level of human control appropriate for each particular AI system and use case, as determined in the requirements definition and risk assessment.

Measure 4: Regulators should require firms to understand their reliance and manage their relationship with third party providers, including monitoring their performance and conducting oversight. To ensure adequate accountability, firms should have a clear service level agreement and contract in place clarifying the scope of the outsourced functions and the responsibility of the service provider. This agreement should contain clear performance indicators and should also clearly determine sanctions for poor performance.

Better define outsourcing type and address separately:

- A distinction should be made between "outsourced functions" and "outsourced AI and ML technology". There is a growing trend toward open source, licensed or bought AI and ML technology. These are two very different outsourced scenarios and should be addressed separately.

Accountability should remain with the firm:

- It is not clear in the Consultation Report whether "testing and oversight of AI and ML" could be outsourced. We recommend accountability over testing and oversight should remain with the firm. Even when some parts of the AI process have been outsourced, or when there is a clear SLA and contract in place. The firm should be held accountable to ensure the same level of quality control internally and with third-party systems.

Applying assurance to the supply chain

- The framework should also recognize and consider how to apply assurance to the supply chain. The System and Organization Controls (SOC) reporting framework was designed for this purpose, especially with respect to securing privacy and data integrity, and should be mentioned here. The Consultation Report relies too heavily on SLAs as a means to manage and provide oversight of outsourced services. But these agreements often do not have right-to-audit clauses or require an audit by the organization. Without the right to audit, an entity may not have the information it needs to ensure its service providers are meeting their obligations. This poses a risk.

Measure 5: Regulators should consider what level of disclosure of the use of AI and ML, is required by firms, including:
(a) Regulators should consider requiring firms to disclose meaningful information to customers and clients around their use of AI and ML that impact client outcomes.
(b) Regulators should consider what type of information they may require from firms using AI and ML to ensure they can have appropriate oversight of those firms.

Strike a balance between transparency and competitive advantage:

- The financial advisor’s fiduciary duty of care and loyalty to the client is the starting point for deciding what to disclose. It is the obligation of the advisor to disclose enough specific information to investors and the user community so they can make informed decisions. However, the advisor must also protect confidentiality around proprietary technologies in order to maintain competitive advantage. Any guidance should try to strike a balance between the two competing needs.

Consider customer and client preferences:

- The measure should include a requirement to facilitate communication back in from customers and clients, including whistle-blowing provisions. This is important because when AI transforms a process/solution it creates a feedback loop and this should be reflected in the measure. Customers and clients should:
  - have a say in the level of information they require,
  - have the ability to dispute an AI-driven decision, and
  - have the option to “opt out” and instead choose a human operator for decision-making, as prescribed by GDPR standards.

Define disclosure type:

- In order to adequately cover all aspects of transparency, this measure should be divided into two categories:
  - data: accuracy, source, confidentiality and privacy, usability (bias or discrimination);
  - models: explainability.

Measure 6: Regulators should consider requiring firms to have appropriate controls in place to ensure that the data that the performance of the ML and AI is dependent on is of sufficient quality to prevent biases and sufficiently broad for a well-founded application AI and ML.

Define bias

- More clarity is needed in defining what is meant by bias. Does the measure relate to inherent biases? Data-driven biases? As well, sometimes bias – for example, a value investing style – is what creates value. Precision is necessary for bias to be addressed adequately.

Data quality can affect bias, but also model performance and robustness:

- It is not clear if this measure implies the need for data quality to reduce bias, or to support general performance and robustness of the model, or both.

Explainability as a way to explore data bias:

- AI explainability can be used to understand and depict the possible bias included in the data.