CFA Society Singapore

146 Robinson Rd, #07-02 UOI Building, Singapore 068909 | Tel: (65) 6323 6679
www.cfasocietysingapore.org E-mail: info@cfasocietysingapore.org

RESPONSE TO CONSULTATION PAPER

Consultation Topic	MAS Consultation Paper on Guidelines on Artificial Intelligence Risk Management
Name / Organisation	**CFA Society Singapore** 1. Abhilash Sivaraman, CFA 2. Chez Anbu, CFA 3. Goh Puay Cheh, CFA 4. Mark Chua, CFA 5. Maurice Teo, CFA 6. Sarah Song, CFA
Email address for any clarification	advocacy@cfasocietysingapore.org
Response to consultation paper	Submission 31 January 2026, via link provided by MAS https://form.gov.sg/690b2a3b024ee5eebbfcf7f1

CFA Society Singapore

146 Robinson Rd, #07-02 UOI Building, Singapore 068909 | Tel: (65) 6323 6679
www.cfasocietysingapore.org E-mail: info@cfasocietysingapore.org

CFA Society Response to MAS Consultation Paper on Guidelines on Artificial Intelligence Risk Management

30 January 2026

Question 1. MAS seeks comments on the application of the Guidelines to all FIs in a proportionate manner, and the guidance on the proportionate application set out in paragraph 1.5 and the Annex of the Guidelines.

We support proportionate application to all FIs, consistent with MAS' risk-based regulatory approach.

- Proportionality should be driven primarily by use-case materiality: (i) impact on customer outcomes, fairness and trust; (ii) complexity; and (iii) reliance/autonomy in decision-making.
- The Guidelines should clearly differentiate AI that makes or materially shapes decisions (e.g., recommendation ranking, underwriting decisions, surveillance triggers) versus AI that assists humans (e.g., summarisation/drafting) – with stricter controls for the former.
- For AI used in regulated activities (financial advice, suitability, portfolio construction, discretionary decision support), MAS could specify a minimum control floor regardless of firm size, to protect clients and avoid uneven standards.
- We recommend MAS include illustrative, tiered examples in the Annex (e.g., low-impact internal productivity tools; medium-impact customer communications; high-impact advisory/portfolio and underwriting use cases) to reduce interpretation variance across firms.

Question 2. MAS seeks comments on the proposed scope of AI use cases, systems and models for the application of the Guidelines.

We broadly agree with the proposed scope covering AI use cases, systems and models, including ML/DL, Generative AI and AI agents.

- We recommend MAS explicitly address AI embedded within third-party tools and enterprise software (e.g., CRM, productivity suites, vendor advisory tooling) – including how Singapore entities should manage risk where models are developed/hosted offshore but used locally.
- Composite advisory architectures are common: a deterministic suitability/product-governance layer combined with probabilistic/GenAI components for explanation or drafting. It would be helpful for MAS to recognise and encourage such constrained AI patterns where appropriate.
- MAS should clarify expectations for distinguishing firm-approved AI use versus employees' personal AI use at work, and reinforce that employment/IT policies should prohibit or tightly restrict unauthorised personal/public AI tools on corporate systems.

Question 3. MAS seeks comments on the proposed responsibilities of the Board and senior management in overseeing AI risk management.

We support MAS' emphasis on Board and senior management oversight, aligned with existing technology risk governance expectations.

- We noted a potential gap in technical understanding at senior levels. Meaningful oversight requires baseline AI competency (training and/or access to expert support) for Board and senior leaders.
- Where AI materially affects customer outcomes (e.g., advisory, suitability, underwriting), senior management should be accountable for calibration, monitoring and escalation processes, including clear decision rights and kill switch authority for severe incidents.
- We recommend at least annual reviews of AI policies, inventories, materiality assessments and controls, with reporting to the Board on high-risk use cases and key risk indicators (KRIs).
- Board reporting for advisory-related AI could include outcome indicators (e.g., suitability breach rates, override/challenge rates, complaint themes, drift/instability metrics, and incident trends).

Question 4. MAS seeks comments on the proposal for FIs to establish a dedicated cross-functional committee to oversee AI risk if the overall AI risk exposure of an FI is deemed material; and how such overall AI risk exposure should be assessed at the organisational level.

We support establishing a dedicated cross-functional committee where overall AI risk exposure is material.

- For high-impact use cases, membership should include technology and data experts, model risk management, compliance/legal, business owners and frontline users (e.g., advisory/wealth, credit) to ensure controls are practical and effective.
- The committee should oversee shadow AI detection and remediation, and ensure that procurement, pre-production approvals and change management gates are consistently applied across business units.
- Overall AI risk exposure should consider both the number/complexity of models and the scale of customer/business exposure (e.g., assets, customer segments, transaction volumes) affected by AI-enabled processes.

Question 5. MAS seeks comments on the proposal for FIs to establish clear definitions, criteria and processes, supported by robust systems, to facilitate the consistent identification of AI across all relevant business and functional areas.

We support the proposal for clear definitions, criteria and processes to identify AI usage across the FI.

- Identification processes should explicitly cover unauthorised AI tools and locally built shadow models (e.g., spreadsheets, scripts, prompt workflows) that can materially affect outcomes without formal governance.
- firms should tag AI by decision role (decisioning vs assistive), level of autonomy/reliance, and whether outputs are deterministic vs probabilistic/generative.

- Controls should also address employee use of personal/public AI tools on corporate devices and networks through IT restrictions and HR policies, to reduce data leakage and compliance blind spots.

Question 6. MAS seeks comments on the proposal for FIs to establish and maintain an accurate and up-to-date inventory of all AI usage.

We support maintaining an accurate and up-to-date AI inventory, and recommend it be impact-focused rather than a mere list of tools.

- Inventories should capture use-case purpose, business owners, affected customer segments/processes, and how AI outputs influence decisions and communications.
- Include materiality ratings (impact, complexity, reliance/autonomy) with rationale; key dependencies (data sources, systems, vendors); and the update mechanism (model/version changes, prompt/template changes, auto-updates).
- For client-facing or outcome-shaping use cases, inventories should link to validation evidence, monitoring dashboards, incident logs and audit trails to support supervisory review.
- Where AI is embedded in third-party services, inventories should capture vendor accountability arrangements, due diligence status, and contingency/exit plans.

Question 7. MAS seeks comments on the proposed risk dimensions of impact, complexity and reliance that should be captured by FIs in AI risk materiality assessments, and whether there are any other risk dimensions that should be included.

We agree that impact, complexity and reliance are strong core dimensions for AI materiality assessments.

- For advisory/suitability contexts, we recommend MAS clarify expectations on reproducibility and auditability: reproducibility at the category/constraint level (e.g., eligible universe, suitability constraints, rationale factors) is a practical standard, even if narrative wording varies.
- We recommend explicit consideration of (i) fairness/non-disadvantage risk; (ii) determinism/stability over time for identical inputs; (iii) client vulnerability and harm severity; (iv) effective reliance (actual override/challenge behaviour); and (v) and third-party opacity (prompt injection, data poisoning, leakage).
- For GenAI in particular, consider truthfulness and misleading-risk (e.g., hallucinations or overconfident framing) as a driver of harm and trust erosion, especially where customers or frontline staff may over-rely on outputs.
- Materiality assessments should also consider whether AI use may affect market integrity and public trust (e.g., misleading content, deepfake or manipulation risks).

Question 8. MAS seeks comments on the proposed standards, processes and controls that should be applied across the entire AI life cycle, and the key areas that FIs should assess for relevance to the AI model, system or use case, and apply in a proportionate manner.

We strongly support the proposed lifecycle controls across inception, development, validation, deployment, monitoring, change management and decommissioning. For high-impact use cases, controls should be calibrated to ensure suitability, auditability and security are not compromised by probabilistic behaviour.

146 Robinson Rd, #07-02 UOI Building, Singapore 068909 | Tel: (65) 6323 6679
www.cfasocietysingapore.org E-mail: info@cfasocietysingapore.org

8.1 Suitability and advisory controls: deterministic constraints with constrained GenAI

- For advisory/portfolio workflows, firms should calibrate AI to be more deterministic where necessary to meet suitability obligations (e.g., hard stops on risk mismatch, concentration, product eligibility).
- GenAI components should be constrained to pre-approved eligible universes and grounded on controlled data sources; narrative variability may be acceptable if the underlying constraints and reasoning factors are reproducible and logged.
- Controls should explicitly address GenAI failure modes such as hallucinations and overconfident framing (e.g., grounding requirements, refusal behaviours, human review for client-facing text, and clear confidence/limitations cues).

8.2 Reproducibility and auditability: decision lineage

- Auditability should include full decision lineage: client inputs used, deterministic rule outcomes, model/version identifiers, prompts/templates, retrieved sources, and timestamps.
- Firms should define practical standards for category-level reproducibility (constraints and rationale factors) and maintain logs sufficient to reconstruct why a recommendation was made and what controls applied.

8.3 Shadow AI, employee controls and data leakage prevention

- Operational controls should proactively detect unauthorised AI usage (shadow AI) and enforce approved tooling via procurement and IT controls.
- Employee use of personal/public AI tools on corporate networks should be restricted through policy, endpoint controls and monitoring, given data confidentiality and leakage risks.

8.4 Third-party AI vendor management and accountability

- Where firms rely on third-party AI, due diligence should assess accuracy, robustness, security, and transparency; contractual clauses should address accountability, incident notification, audit rights and change management.
- MAS could clarify how AI vendor management should integrate with existing TRM and outsourcing guidelines, including expectations where liability transfer may be limited in practice.

8.5 Human oversight and over-reliance risks

- Advisors remain ultimately responsible for recommendations except in the cases of AI based robo-advisories; governance should include override/challenge expectations and monitoring for automation bias (e.g., unusually low challenge rates).
- Firms should design UI and workflows to encourage critical review, and establish escalation routes where AI outputs conflict with policy or client interests.

8.6 Transparency and client disclosure

- Where AI materially influences advice or customer outcomes, disclosure to clients regarding AI's role and limitations can improve transparency and market-driven accountability, while preserving the FI's responsibility for outcomes.

Question 9. MAS seeks comments on any aspects of the Guidelines that have not been covered in earlier questions, as well as aspects of AI risk management that have not been covered in the proposed Guidelines.

Overall, the Guidelines are comprehensive. We recommend additional clarifications in three areas:

- Advisory and suitability: include a short consolidated sub-section (or Annex examples) for AI in financial advisory/portfolio management, covering suitability constraints, category-level reproducibility expectations, and audit lineage.
- Third-party embedded AI and role clarity: clarify expectations for AI embedded in vendor tools and the respective responsibilities of manufacturers versus distributors of AI-enabled services in financial contexts.
- Vendor accountability and outsourcing alignment: clarify how accountability, due diligence and ongoing monitoring expectations interact with outsourcing requirements and cross-border group governance.
- IP and innovation: maintain flexibility and avoid approaches that force disclosure of proprietary methods; focus on governance, monitoring and outcomes rather than requiring detailed model disclosure.
- Regulatory evolution: consider a periodic review mechanism (e.g., half-yearly or annual) to update expectations as AI transparency and stability improve, potentially shifting from prescriptive controls to more principle-based guidance over time.
- Cross-jurisdictional alignment: benchmark with international regulators where appropriate to reduce fragmentation and implementation burden for globally operating firms.

Question 10. MAS seeks comments on the proposed transition period of 12 months.

We agree that a 12-month transition period is reasonable, provided implementation is prioritised by materiality and client impact.

- We suggest MAS encourage interim milestones for high-impact use cases (e.g., within the first 3–6 months: complete AI identification and impact-focused inventory; classify use cases by materiality; implement minimum safeguards for advisory/suitability and other outcome-shaping uses).
- MAS could also signal an intention to re-evaluate the Guidelines on a scheduled basis (e.g., within 6–12 months after issuance) to incorporate lessons learned and evolving risks, particularly for GenAI and agents.