Principles for responsible use

Dr Alexander Ressos, TEP
Of Counsel, KGP Legal LLC
Attorney-at-Law (Switzerland and Germany), Solicitor (England & Wales and Hong Kong)

The FEAT Principles of 2018

Singapore was among the first jurisdictions to issue principles specifically for the use of artificial intelligence in finance. On 12 November 2018 MAS published the Principles to Promote Fairness, Ethics, Accountability and Transparency in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector, known by their initials as FEAT. A committee co-chaired by David Hardoon, then MAS Chief Data Officer, and Hsieh Fu Hua of PrimePartners drafted them with financial institutions, industry associations, fintech firms, technology providers and academics.[1]

The fourteen principles are brief. Under fairness, individuals or groups are not to be systematically disadvantaged through AI-driven decisions unless the decisions can be justified, while the use of personal attributes as input factors must itself be justified. Data and models are to be reviewed regularly for accuracy, relevance and unintentional bias. Under ethics, AI-driven decisions are held to at least the same ethical standards as decisions taken by humans. Under accountability, the use of artificial intelligence must be approved by the appropriate internal authority. Firms are accountable for both internally developed and externally sourced models and are expected to raise the awareness of management and the board. Data subjects should have channels to enquire about, appeal and request review of decisions that affect them. Under transparency, the use of such systems is to be disclosed proactively, with clear explanations available to data subjects on request about the data used and the consequences of the decision.[2]

The principles were deliberately high-level and non-binding. Their significance lay in the method. MAS chose to establish a shared vocabulary with industry before it set supervisory expectations, so that the later guidelines could build on concepts that the institutions had helped to define.

Veritas: from principles to methods

Principles alone do not tell a bank how to measure whether a credit model is fair. On 13 November 2019 MAS announced Veritas, an initiative with a first consortium of seventeen members to develop open-source tools that would allow financial institutions to assess their systems against the FEAT Principles, beginning with customer marketing, risk scoring and fraud detection.[3]

The work proceeded in three phases. The first, which ended on 6 January 2021, produced fairness metrics for credit risk scoring and customer marketing and published the code. The second, concluded on 4 February 2022 with twenty-seven members, developed a full assessment methodology for fairness, ethics, accountability and transparency together with version 1 of the Veritas Toolkit. The

Copyright © 2026 Dr Alexander Ressos, TEP – Of Counsel, KGP Legal LLC. SingaporeIQ is an information website and gives no legal advice. For Singapore law advice through KGP Legal LLC, please write to alexander.ressos@kgplegal.com.sg

third phase grew to thirty-two members, with newcomers that included the reinsurer Munich Re, Visa and the Singapore Exchange. Seven financial institutions piloted the integration of the methodology into their governance frameworks before the phase ended on 26 June 2023 with version 2.0 of the Toolkit.[4]

Veritas made the FEAT Principles operational for conventional machine learning. It did not anticipate generative AI, whose outputs are open-ended text, which makes the fairness metrics developed for credit scores difficult to apply.

The national framework and data protection

The financial sector rules sit within a national framework. The Infocomm Media Development Authority and the Personal Data Protection Commission (PDPC) published a Model AI Governance Framework in January 2019 and a second edition in January 2020, followed by a framework for generative AI in 2024.[5] IMDA has since added a Model AI Governance Framework for Agentic AI, whose revised version 1.5 appeared on 20 May 2026.[6] MAS refers financial institutions to the agentic framework as a practical reference for the management of agentic and multi-agent systems and intends to consult the sector in 2027 on additional guidance of its own.[7]

Data protection law applies to every model trained or operated on personal data. The PDPC issued Advisory Guidelines on the Use of Personal Data in AI Recommendation and Decision Systems on 1 March 2024. They explain that an organisation may use personal data without consent to develop or enhance a product, or to improve operational efficiency, under the business improvement exception of the Personal Data Protection Act 2012. A separate research exception covers commercial research without a product development roadmap. For banks, the practical effect is that historical customer data can in principle be used to train internal models. They must still respect the consent and notification obligations when the system is deployed and the accountability obligation throughout.[8] The banking secrecy rules of the Banking Act, examined in Part E, remain applicable in addition.

Project MindForge

The arrival of generative AI prompted a second industry project. In its first phase, Project MindForge brought together six banks, the Association of Banks in Singapore and three technology and consulting partners. Its first phase produced a risk framework for generative AI, a platform-agnostic reference architecture and a white paper on the emerging risks and opportunities for banks, whose executive summary appeared on 15 November 2023.[9]

The second phase widened the project to insurance and asset management and from generative AI to conventional and agentic AI. Thirty-four participants took part, among them twelve banks, eight insurers, four capital market firms, five industry associations and five technology and consulting partners. They produced an AI Risk Management Toolkit consisting of an executive handbook, an operationalisation handbook and a collection of implementation examples, designed as companions to MAS's guidelines.[10] The pattern is the same as with Veritas. Industry and regulator work out the practice together before supervisory expectations follow.

Copyright © 2026 Dr Alexander Ressos, TEP – Of Counsel, KGP Legal LLC. SingaporeIQ is an information website and gives no legal advice. For Singapore law advice through KGP Legal LLC, please write to alexander.ressos@kgplegal.com.sg

The Guidelines on Artificial Intelligence Risk Management

The step from good practice to supervisory expectation came in two stages. MAS first reviewed the artificial intelligence model risk management of selected banks in mid-2024 and published its observations in December 2024, announcing that it was considering supervisory guidance for all financial institutions in the following year.[11] On 13 November 2025 it issued a consultation paper on Guidelines on Artificial Intelligence Risk Management, with comments due by 31 January 2026. Ho Hern Shin, Deputy Managing Director, described them as "proportionate, risk-based guidelines" intended to enable responsible innovation.[12]

MAS published its response on 7 October 2026 together with the final Guidelines, which take effect on 7 October 2027.[13] The list of respondents runs to ninety-five named or identified entries, to which come forty respondents who asked for their identity to remain confidential. They include industry associations, banks, insurers, technology firms and individuals.[14] The final text retains the architecture of the consultation and makes several adjustments in favour of proportionality.

The scope is wide. The Guidelines apply to all financial institutions, wherever incorporated, and on a group basis for institutions under consolidated supervision or owning critical information infrastructure. Artificial intelligence is defined as machine-based systems or models that derive outputs such as estimates, predictions, content, summaries, recommendations or decisions through learned premises. Tools whose outputs rest solely on predefined rules are excluded. Logistic regression and gradient boosting, deep learning, natural language processing, computer vision, generative AI and AI agents fall inside the definition. Rule-based investment formulas, expert systems built on if-then rules, traditional robotic process automation and Monte Carlo simulations not derived from training data fall outside it. Third-party AI is included, even where it is embedded in a service that is not marketed as artificial intelligence.[15]

Proportionality operates at two levels. An institution whose use is limited to assistive tools, where poor performance or unavailability is unlikely to have a material adverse impact, may apply basic governance policies only. Examples are the drafting of e-mails, the summarising of documents and meeting notes and the generation of charts for internal reference. All other institutions must implement the full set of expectations in a risk-proportionate manner. At the level of the individual use case, each application must be identified, recorded in an inventory and assessed for risk materiality along at least three dimensions: impact, complexity and reliance. MAS retained complexity despite objections and folded explainability and third-party considerations into the definitions. Reliance now focuses on the autonomy granted to the system, the degree of human oversight and the dependence on third parties.[16]

The lifecycle controls cover data management, fairness, transparency and explainability, human oversight, the selection of models, evaluation and testing, technology and cybersecurity, reproducibility and auditability, pre-deployment reviews, post-deployment monitoring and change management. The board approves key frameworks and the risk appetite, while senior management implements them. MAS added that the board need not possess hands-on technical expertise, although it must be able to challenge management effectively. The timetable was extended after many

Copyright © 2026 Dr Alexander Ressos, TEP – Of Counsel, KGP Legal LLC. SingaporeIQ is an information website and gives no legal advice. For Singapore law advice through KGP Legal LLC, please write to alexander.ressos@kgplegal.com.sg

respondents asked for eighteen to twenty-four months. Institutions must establish their oversight frameworks, inventories and risk materiality assessments within twelve months, and complete the lifecycle controls and capability requirements within twenty-four months. High-risk use cases should receive appropriate controls as soon as possible.[17]

Common misconception

It is often said that Singapore regulates artificial intelligence in finance through voluntary principles only. That description no longer fits. The Guidelines on Artificial Intelligence Risk Management of October 2026 set supervisory expectations for every financial institution. They carry no criminal penalties of their own. MAS assesses compliance in its supervision and may draw on its general powers where an institution's risk management is inadequate.

Notes

1. MAS, "MAS introduces new FEAT Principles to promote responsible use of AI and data analytics", media release, 12 November 2018: FEAT Committee co-chaired by Dr David Hardoon, MAS Chief Data Officer, and Mr Hsieh Fu Hua, Co-Founder and Advisor, PrimePartners; feedback from FIs, industry associations, FinTech firms, technology providers and academia.
2. MAS, *Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector* (version updated 7 February 2019, mas.gov.sg), principles 1–14: approval "by the appropriate internal authority"; firms "accountable for both internally developed and externally sourced AIDA models"; channels to "enquire about, submit appeals for and request reviews of AIDA-driven decisions"; proactive disclosure and explanations on request.
3. MAS, "MAS partners financial industry to create framework for responsible use of AI", media release, 13 November 2019: Veritas consortium of MAS, SGInnovate, EY and 14 financial institutions; initial use cases in customer marketing, risk scoring and fraud detection.
4. MAS, "Veritas Initiative" (mas.gov.sg, accessed November 2026): Phase 1 with 25 members concluded on 6 January 2021; Phase 2 with 27 members concluded on 4 February 2022 with five white papers and Toolkit version 1; Phase 3 with 32 members concluded on 26 June 2023 with Toolkit version 2.0; seven FIs (BNY Mellon, DBS, HSBC, OCBC, Singlife, Standard Chartered and UOB) piloted integration with their governance frameworks; new Phase 3 members included Munich Re, SGX, Singlife, the Singapore FinTech Association and Visa; the Toolkit is open source on GitHub.
5. IMDA and PDPC, *Model Artificial Intelligence Governance Framework*, first edition January 2019, second edition January 2020; *Model AI Governance Framework for Generative AI*, May 2024.
6. IMDA, *Model AI Governance Framework for Agentic AI*, Version 1.5, published 20 May 2026, updated 5 June 2026 (imda.gov.sg, accessed November 2026), "What's new in this version": incorporates feedback "from 60+ companies since v1.0".
7. MAS, "MAS sets out supervisory expectations on responsible AI adoption by financial institutions", media release, 7 October 2026 (mas.gov.sg): consultation on additional agentic AI guidance planned for 2027 – MAS, Response to Feedback, 7 October 2026, para 12.9: MAS has included "the national-level Agentic AI Model Governance Framework published by the Infocomm Media Development Authority as a practical reference"; MAS intends to consult the sector separately on additional guidance specific to agentic AI.

Copyright © 2026 Dr Alexander Ressos, TEP – Of Counsel, KGP Legal LLC. SingaporeIQ is an information website and gives no legal advice. For Singapore law advice through KGP Legal LLC, please write to alexander.ressos@kgplegal.com.sg

8. PDPC, *Advisory Guidelines on Use of Personal Data in AI Recommendation and Decision Systems*, issued 1 March 2024, paras 4.2 and 5.1, with references to Part 5 of the First Schedule and Divisions 2 and 3 of Part 2 of the Second Schedule to the PDPA – PDPC, *Advisory Guidelines*, 1 March 2024, Part IV, sections 9 ("Consent and Notification Obligations") and 10 ("The Accountability Obligation").
9. MAS, "Project MindForge" (mas.gov.sg, page dated 13 November 2025, accessed November 2026): Phase 1 participants DBS, OCBC, UOB, Standard Chartered, Citi Singapore, HSBC, ABS, Accenture, Google Cloud and Microsoft; outputs: generative AI risk framework, reference architecture and the white paper *Emerging Risks and Opportunities of Generative AI for Banks*; executive summary released on 15 November 2023.
10. MAS, "Project MindForge" (above): Phase 2 with 34 participants (12 banks, 8 insurers, 4 capital market firms, 5 industry associations, 5 technology and consulting partners); AI Risk Management Toolkit comprising an Executive Handbook, an Operationalisation Handbook and Implementation Examples, to be updated in line with relevant MAS guidelines.
11. MAS, "Artificial Intelligence Model Risk Management" (mas.gov.sg, published 5 December 2024): "MAS conducted a thematic review of banks' AI (including Generative AI) model risk management practices in mid-2024"; Information Paper, December 2024, para 8.2: "MAS is also considering supervisory guidance for all FIs next year".
12. MAS, media release on the consultation on Guidelines for Artificial Intelligence Risk Management, 13 November 2025 (mas.gov.sg): consultation until 31 January 2026; application to all FIs; coverage of generative AI and AI agents; quotation of Ms Ho Hern Shin, Deputy Managing Director.
13. MAS, *Guidelines on Artificial Intelligence Risk Management*, 7 October 2026, para 1.8: the Guidelines "take effect on 7 October 2027"; Sections 3 and 4 from 7 October 2027, Sections 5 and 6 by 7 October 2028.
14. MAS, Response to Feedback, 7 October 2026, para 1.3 (final Guidelines accessible through the response) and Annex A (95 numbered respondents, plus 11 and 29 respondents who requested confidentiality of identity, respectively of identity and submission).
15. MAS, Response to Feedback, 7 October 2026, paras 2.3 and 2.7 – MAS, Response to Feedback, 7 October 2026, paras 2.8–2.11, including the table of illustrative examples within and outside the scope of the AI definition.
16. MAS, Response to Feedback, 7 October 2026, paras 3.6–3.9: revised criterion based on whether poor performance or unavailability "is unlikely to have a material adverse impact on the FI, its customers or other stakeholders"; examples of assistive use in para 2.4 of the Guidelines – MAS, Response to Feedback, 7 October 2026, paras 7.2, 7.5–7.9 and 7.13.
17. MAS, Response to Feedback, 7 October 2026, paras 4.7, 4.8, 4.10(a) and 10.1 – MAS, Response to Feedback, 7 October 2026, paras 13.1–13.6: 12 months for Sections 3 and 4 of the Guidelines and for basic AI governance; 24 months for full implementation including Sections 5 and 6; high-risk use cases not to await the end of the transition period.

Copyright © 2026 Dr Alexander Ressos, TEP – Of Counsel, KGP Legal LLC. SingaporeIQ is an information website and gives no legal advice. For Singapore law advice through KGP Legal LLC, please write to alexander.ressos@kgplegal.com.sg