

INTERNATIONAL AI IN FINANCIAL SERVICE REVIEW

2026/27
EDITION



Beaumont Capital Markets

www.beaumont-capitalmarkets.co.uk

JAPAN

JAPAN'S TECHNOLOGY-NEUTRAL APPROACH TO CUSTOMER-FACING GENERATIVE AI IN FINANCE

AMT



Takashi Nakazaki

Partner



takashi.nakazaki_grp@amt-law.com



+81-3-6775-1086



www.amt-law.com



BIO

Takashi Nakazaki is a Partner at Anderson Mori & Tomotsune. His practice focuses on data protection and privacy, AI, information security, intellectual property, software and licensing, and payment services, together with related technology, e-commerce and corporate matters.

He is recognised in leading legal rankings and has published Japanese-language books on the GDPR (2018), the metaverse (2023) and generative AI (2024). He has served on expert committees for METI, MIC, the Cabinet Office and other bodies, including the working group that drafted the AI Guidelines for Business, and has led research commissioned by the Personal Information Protection Commission. He has been an auditor of the Japan Institute for Health Security since April 2025.

Mr Nakazaki co-founded the IAPP Tokyo KnowledgeNet Chapter, serves on the IAPP Asia Advisory Board, and teaches AI law at Keio University Law School. He is Associations and Committees Liaison Officer of the IBA Technology Law Committee and a member of AIPPI's Standing Committee on AI, IT and Internet.

He received his LL.B. from the University of Tokyo and his LL.M. from Columbia Law School, and previously worked at Arnold & Porter LLP.

AMT



Tomoaki Katayama

Partner



tomoaki.katayama@amt-law.com



+81-3-6775-1590



www.amt-law.com



BIO

Tomoaki Katayama is a partner at Anderson Mori & Tomotsune. Prior to joining the firm in February 2020, he worked at Japan Post Bank Co., Ltd. (April 2013–August 2013), Taisei Corporation (September 2013–April 2017), Sato Sogo Law Office (May 2017–October 2017), and So & Sato Law Offices (November 2017–January 2020). Drawing on the broad perspective he has developed through these diverse professional experiences, he advises on FinTech and financial regulation, with deep and extensive expertise in digital assets, including crypto-assets, stablecoins, and NFTs, as well as related cutting-edge technologies such as blockchain technology and its applications, including DeFi. He also handles matters relating to personal data protection, e-commerce, financial transactions, and gaming law.

Publications

6500 Legal Countermeasures for Financial Institutions (February 2026)
Global Legal Insights - Fintech 2025 Japan Chapter (September 2025)
Chambers Global Practice Guides' on Blockchain 2025 - Law & Practice (June 2025)

AMT



Takeshi Nagase

Partner



takeshi.nagase@amt-law.com



+81-3-6775-1200



www.amt-law.com



BIO

Takeshi Nagase is a partner at Anderson Mori & Tomotsune. Between 2013 and 2014, Takeshi served on secondment in the disclosure department at the Financial Services Agency of Japan (FSA). Additionally, he handled a broad range of finance and corporate transactions on a secondment stint with the legal department of a major Japanese securities firm from 2015 to 2017. As a result of the unique perspective he has gained from these professional experiences, Takeshi has extended his focus to crypto-asset laws, including regulatory requirements applicable to the registration of crypto-asset exchange service providers, initial coin offerings, stablecoin businesses, and NFT businesses.

AMT



Daisuke Iguchi

Special Counsel



daisuke.iguchi@amt-law.com



+81-3-6775-1885



www.amt-law.com



BIO

Daisuke Iguchi is a Special Counsel at Anderson Mori & Tomotsune. He has gained experience through a secondment to a bank establishment preparatory company, where he was engaged in various legal matters related to the establishment of a bank, and as a fixed-term government employee at the Financial Services Agency, where he handled various legal inquiries under the Payment Services Act and the Money Lending Business Act.

Drawing on this experience, he specializes in cross-sectional analysis of various financial regulations, including the Banking Act, the Payment Services Act, the Money Lending Business Act, the Installment Sales Act, and the Financial Instruments and Exchange Act, as well as support in dealing with the authorities.

JAPAN'S TECHNOLOGY-NEUTRAL APPROACH TO CUSTOMER-FACING GENERATIVE AI IN FINANCE

1. Introduction: The Shift Towards Customer-Facing Generative AI

In Japan's financial sector, generative AI ("GenAI") is beginning to move beyond internal efficiency tools towards customer-facing services. When the Financial Services Agency of Japan

(the "FSA") surveyed 130 financial institutions between October and November 2024, direct use of GenAI in customer-facing services was still rare. By contrast, the AI Discussion Paper (Version 1.1), "Preliminary Discussion Points for Promoting the Sound Utilization of AI in the Financial Sector" (the "AIDP"), published on 3 March 2026, records that institutions have since begun launching, or concretely considering, customer-facing

GenAI services within limited scopes and conditions.

"AI agents" have also become a realistic object of consideration. Unlike conventional chatbots, these systems may access multiple systems and data sources and autonomously perform defined tasks. GenAI is therefore evolving from an employee-support tool into a technology capable of forming part of the process through which financial services are delivered.¹

Japan's regulatory response is distinctive. Rather than establishing a comprehensive financial-sector regime specifically for GenAI, Japan applies existing financial regulation in a technology-neutral manner, while addressing uncertainty through risk-based governance and public-private dialogue. The AIDP also highlights the "risk of inaction": the possibility that overly cautious institutions may fall behind technological innovation and become less able to provide high-quality financial services over the medium to long term. This article examines that approach, focusing on customer-facing GenAI, securities regulation, personal data protection and key risk-management issues.



2. Japan's AI Governance: A Promotion-Oriented Framework Act and Soft Law

Japan's AI policy seeks to reconcile innovation with risk management. The interim report of the government's AI Institution Study Group similarly identifies "balancing innovation promotion and risk response" as a central policy objective, favouring a combination of legislation and soft law while respecting business autonomy.²

This approach is reflected in the Act on the Promotion of Research and Development and Utilization of Artificial Intelligence-Related Technologies (Act No. 53 of 2025; the "Japan AI Act"), promulgated on 4 June 2025. Article 3(2) provides that, given AI's potential to improve efficiency, enhance business activities and create new industries, its research, development and utilization should be promoted to maintain domestic capabilities and strengthen international competitiveness. Article 3(4), however, recognizes that improper or inappropriate use may facilitate crime, leakage of personal information, copyright infringement and other harms, and calls for measures to ensure proper implementation.

The AI Act therefore differs in character from the EU AI Act. It does not establish a comprehensive risk-classification regime with detailed obligations and sanctions. Instead, it provides a framework for promoting AI-related measures comprehensively and systematically (Article 1), requires an AI Basic Plan (Article 18), and establishes the AI Strategy Headquarters within the Cabinet (Articles 19 and 20). The first plan was adopted by Cabinet decision on 23 December 2025, and the Phase II Artificial Intelligence Basic Plan was adopted by Cabinet decision on 14 July 2026. At the same time, it is not merely a promotion statute: Article 13 contemplates guidelines to ensure proper implementation, while Article 16 provides for analysis of cases involving improper AI use and measures including guidance and advice. It is therefore best characterized as a promotion-oriented and risk-responsive framework act.

1. See Ministry of Internal Affairs and Communications and Ministry of Economy, Trade and Industry, AI Guidelines for Business, Version 1.2 (31 March 2026), defining an AI agent as an AI system that senses its environment and acts autonomously to achieve a specified goal (author's translation).
2. Cabinet Office, AI Strategy Council / AI Institution Study Group: Interim Report (February 2025) (in Japanese; translations of quoted language are the author's).

Soft law remains important. The Social Principles of Human-Centric AI set out overarching societal principles. Separately, the AI Guidelines for Business provide practical guidance for actors engaged in AI development, provision and use; Version 1.2 was published on 31 March 2026.³ Japan's overall model accordingly combines a cross-sector framework act, soft law and existing sector-specific regulation.

This structure is particularly clear in finance. The Banking Act, the Financial Instruments and Exchange Act ("FIEA"), the Insurance Business Act, the Act on the Protection of Personal Information ("APPI") and other existing laws apply regardless of whether AI is used. The FSA does not exclude legislation where a material regulatory gap emerges, but has indicated that interpretation or revision of existing principles and guidelines should generally be considered first.

Japan's Copyright Act is also relevant to AI development. Article 30-4 permits certain uses of works for information analysis and other purposes that do not involve enjoying the thoughts or sentiments expressed in the work, and AI training may fall within its scope. The exception is not unlimited, particularly where a purpose of enjoyment is also present or the use would unreasonably prejudice the interests of the copyright owner.

3. The Use of Generative AI by Japanese Financial Institutions

In the survey underlying the AIDP, more than 90% of responding institutions used conventional AI or GenAI in some form. GenAI use had expanded rapidly in internal operations, including document summarization, translation, proofreading, internal FAQs and system development, and more than 70% of institutions using AI in some form permitted broad GenAI use by general staff.

The AIDP broadly distinguishes three categories. The first is internal use, such as drafting, translation and information retrieval. The second is use supporting customer services or customer-affecting operations through a human intermediary, such as proposed responses for call-centre staff or drafts used in decision-making. The third is direct customer use, where GenAI output itself is presented to customers.

At the time of the 2024 survey, the third category was extremely limited. Version 1.1, however, records that direct customer use has begun in limited areas such as enquiry handling and administrative procedures. More recent Bank of Japan data reinforce this picture. In its FY2022026 survey of 150 financial institutions, 79.3% reported that they were already using GenAI and a further 11.3% were trailing it. Use had expanded from general administrative tasks into core operations involving customer information, although only a small number of institutions were directly presenting GenAI-generated output to customers.⁴

The risks change accordingly. Internal use primarily raises issues of confidentiality, erroneous output, copyright and employee misuse. Direct customer use may additionally trigger financial-regulatory concerns, including solicitation, suitability, customer disclosure and the effects of misinformation. If AI agents begin autonomously executing transactions or operational processes, questions of authority, accountability and system risk will become more significant.

“

Japan's overall model accordingly combines a cross-sector framework act, soft law and existing sector-specific regulation.

4. The FSA's Approach: Risk-Based Lifecycle Management

The AIDP is not a supervisory guideline and does not itself impose new legal obligations. It sets out preliminary issues for dialogue with the industry and the FSA's policy direction. Its basic approach rests on technology neutrality and risk-based regulation.

For customer-facing GenAI, the AIDP is especially notable for organizing risk-mitigation practices into a lifecycle framework. First, design and pre-deployment verification. Institutions may combine retrieval-augmented generation ("RAG"), system prompts, model selection, fine-tuning and output filters as layered guardrails. Deployment may begin with relatively low-risk services and expand as experience accumulates, with testing conducted before launch.

3. Ministry of Internal Affairs and Communications and Ministry of Economy, Trade and Industry, AI Guidelines for Business (Version 1.2) (31 March 2026).

4. Bank of Japan, Use and Risk Management of Generative AI by Japanese Financial Institutions -Based on the Results of FY2026 Survey- (24 August 2026).

Second, customer explanation and warnings. Among the practices discussed in the AIDP are informing customers that they are interacting with GenAI, warning them that responses may contain errors, identifying underlying information sources where appropriate, and providing a route to human support.⁵

“

The AIDP also refers to “agile governance”: repeatedly reassessing risks, objectives and controls rather than treating governance as a one-off exercise.

Third, continuous verification and monitoring. Relevant controls include retaining and reviewing conversation logs, testing for skewed product recommendations, documenting recommendation logic and, where appropriate, obtaining third-party review.

Fourth, firm-wide governance. Senior-management involvement, clear responsibility for AI, employee literacy, risk-proportionate

controls and continuous review across the AI lifecycle are important. The AIDP also refers to “agile governance”: repeatedly reassessing risks, objectives and controls rather than treating governance as a one-off exercise.

These measures are not legal duties in themselves. Nevertheless, the FSA has indicated that it will refer to the practices described in the AIDP when responding to consultations concerning AI. They are therefore likely to become important practical reference points for institutions introducing customer-facing GenAI.

5. Securities Regulation: Applying Existing Rules to Generative AI

One important feature of the AIDP is that it addresses concrete regulatory questions facing financial instruments business operators rather than merely stating high-level AI principles.

The first concerns intra-group AI development and non-public information. Article 44-3(1)(iv) of the FIEA, together with Article 153(1)(vii) of the Cabinet Office Order on Financial Instruments Business, etc. (the “**FIB Order**”), restricts certain exchanges of non-public issuer information between a Type I financial instruments business operator engaged in securities-related business and its parent or subsidiary corporations.

Article 153(1)(vii)(g), however, provides an exception for information necessary for the maintenance and management of an electronic data processing system where appropriate safeguards against leakage are in place. On the premise that “maintenance and management” has historically been understood to include system development, the AIDP indicates that non-public information need not necessarily be removed in advance from customer conversation data used for AI-system development. AI development therefore does not automatically attract stricter information-sharing rules than conventional system development.

The second issue is access to corporate-related information. Article 40(ii) of the FIEA and Article 123(1)(v) of the FIB Order require necessary and appropriate measures to prevent unfair trading in connection with the management of corporate-related information. Japan Securities Dealers Association rules also require controls preventing such information from reaching divisions that do not need it.

The AIDP indicates that granting data scientists or similar personnel access does not, by itself, establish a control failure. The analysis remains substantive: institutions should identify the purpose and manner of AI use, access rights and information-management procedures, and assess whether the overall control framework is appropriate.

The third issue is AI product recommendations and “solicitation”. The FIEA contains no general statutory definition of solicitation. The AIDP states that solicitation is generally understood as conduct directed at specific users with the aim of inducing financial transactions, and that the same concept applies where AI is used.

5. FSA, Preliminary Discussion Points for Promoting the Sound Utilization of AI in the Financial Sector (Version 1.1) (3 March 2026) (Japanese). The FSA's English webpage announces Version 1.1, but a full English translation of Version 1.1 was not located as of 29 August 2026; translations of Version 1.1-specific passages in this article are the author's.

Where GenAI activity constitutes solicitation, existing conduct rules apply. Article 40(i) of the FIEA contains the suitability principle, while Article 38 prohibits, among other things, false statements and solicitation through conclusive evaluations of uncertain matters. The central regulatory question is therefore not whether a human or AI communicated with the customer, but whether the AI interaction functioned to induce that customer to enter into a financial transaction.

The FSA has indicated that, as AI-specific issues develop, relevant considerations may include AI's influence on investor decision-making and the degree of human involvement. This illustrates Japan's broader model: existing law is applied first, with ambiguities clarified through regulatory dialogue.

The FinTech Support Desk and FinTech Proof-of-Concept Hub support that process. Where a formal written interpretation is required, mechanisms such as the Grey Zone Elimination System under the Act on Strengthening Industrial Competitiveness may also be used.

6. Personal Data Protection: Key Issues for Generative AI

Personal data protection is among the most significant legal challenges for Japanese financial institutions using GenAI.

2026 Amendment to the APPI. As of August 2026, Japan has also enacted a significant amendment to the APPI. The amendment was passed by the Diet on July 10, 2026 and promulgated on July 17, 2026. Among other changes, it creates exceptions to certain consent requirements where personal data is used exclusively for the creation of statistical information or similar outputs, a category that the PPC expressly states may include certain forms of AI development. The amendment is intended in part to facilitate data use, including for AI, while maintaining safeguards against uses that could adversely affect individuals' rights and interests. However, most provisions of the amendment have not yet entered into force and will take effect on a date to be specified by Cabinet Order within two years of promulgation. Accordingly, financial institutions should distinguish between the APPI requirements currently in force and the forthcoming framework under the 2026 amendment.

The AIDP discusses several recurring APPI issues by reference to views expressed by experts in the AI Public-Private Forum. They should therefore not all be read as new, definitive interpretations by the FSA or the Personal Information Protection Commission.

The first is the purpose of use for AI development and training. Under the APPI, personal information must generally be handled within a specified purpose of use. Existing guidance recognizes that processing undertaken solely to produce statistical or otherwise non-identifiable results need not always be separately stated as an independent purpose.⁶ The AIDP records expert views that similar reasoning may be relevant to certain forms of AI development and training. Where AI is used to make determinations about, or provide information to, individual customers, however, institutions should consider whether such handling can reasonably be anticipated by the data subject.

The second issue is the input of personal data into external GenAI services. If the AI provider processes personal data only within the institution's instructions, the arrangement may constitute an entrustment of personal-data handling.

The institution must then supervise the provider and confirm that input data is not used outside the institution's instructions for additional model training or the provider's own purposes. Enterprise contracts should therefore address data-use scope, model training, retention, sub-contracting, security and incident reporting.

The third issue is outsourced AI development, where similar controls over purpose, access, security and supervision apply.

“

The central regulatory question is therefore not whether a human or AI communicated with the customer, but whether the AI interaction functioned to induce that customer to enter into a financial transaction.

6. Personal Information Protection Commission, Q&A on the Guidelines under the Act on the Protection of Personal Information (last revised 1 July 2025), 2-5.

The fourth is cross-border use. Whether the APPI's foreign-third-party rules apply is generally analyzed by reference to the location of the receiving business operator; an overseas server location alone does not necessarily make a transfer a foreign-third-party provision. Server location nevertheless remains relevant to the "external environment" that must be understood as part of security control measures.

“

Hallucination is likewise better treated as a residual risk to be managed through RAG, source display, guardrails, human confirmation and escalation than as a risk capable of complete elimination.

In practice, key safeguards include restricting unnecessary input of personal data into GenAI, using services contractually and technically restricted from using customer inputs for additional model training or other provider purposes.

7. Explainability, Hallucination and Model Risk

As GenAI becomes customer-facing, reliability and explainability become increasingly important. Where AI is used in significant financial judgments, institutions should be able to assess the reasonableness of outcomes and provide an appropriate explanation where necessary.

Because the internal reasoning behind individual GenAI outputs may be difficult to explain fully, governance often shifts towards process-level verifiability: recording the data used, RAG sources, prompts, outputs and human review. Hallucination is likewise better treated as a residual risk to be managed through RAG, source display, guardrails, human confirmation and escalation than as a risk capable of complete elimination.

GenAI also challenges conventional model-risk management. Institutions are seeking to adapt existing model-risk frameworks,⁷ but GenAI may produce different outputs from identical inputs, evaluation methodologies remain immature and external foundation models may change frequently. Maintaining an inventory of significant AI systems, including their purposes, assumptions, limitations and vendor changes, is therefore increasingly important.

8. Third Parties and Security

Financial institutions' GenAI use depends heavily on external foundation models and cloud services, making AI risk closely connected to third-party risk.

Existing outsourcing frameworks remain relevant, but AI providers raise additional issues concerning data use, model changes, sub-contracting, service continuity, security, intellectual property and vendor lock-in. Concentration on a small number of foundation-model providers may also create common vulnerabilities across the financial system.⁸

Cybersecurity risks include not only confidential information entered into prompts, but also prompt injection, data poisoning, model extraction and unauthorised employee use of external AI services. At the same time, GenAI may be abused for sophisticated phishing, deepfake-enabled identity fraud and forged documentation, creating new challenges for AML/CFT and fraud controls.

As AI agents increasingly perform transactions or fund transfers autonomously, entitlement management, authentication, human oversight and mechanisms for suspension or intervention will become increasingly important.

7. FSA, Principles for Model Risk Management (November 2021); see also FSA, progress report on advancing model risk management at financial institutions (December 2024).

8. Financial Stability Board, The Financial Stability Implications of Artificial Intelligence (November 2024).

9. Conclusion

Japan has chosen not to begin with detailed finance-specific GenAI rules. Instead, it applies existing regulation in a technology-neutral manner and develops regulatory practice through risk-based governance and continuous public-private dialogue.

The AIDP shows that this model is moving beyond abstract principles into concrete questions concerning customer-facing GenAI, securities regulation, personal data protection and AI governance. Its strength is flexibility: Japan can respond to fast-moving technology without unnecessarily duplicating existing financial regulation. Its weakness is that institutions may still lack a sufficiently clear benchmark for “how much is enough”.⁹

The ultimate test is therefore not whether Japan eventually enacts AI-specific financial regulation. It is whether existing technology-neutral rules can be translated, through guidance, interpretation and supervisory dialogue, into sufficiently clear and practical standards as GenAI and AI agents move closer to the core of financial services.

The expansion of customer-facing GenAI and agentic systems in 2026 and 2027 will provide a significant test of that iterative regulatory model.

“

Its weakness is that institutions may still lack a sufficiently clear benchmark for “how much is enough”.

9. See Daiwa Institute of Research, The Regulatory Approach to AI in the Financial Sector: Key Points of the AI Discussion Paper (Version 1.1) and International Comparison (15 April 2026) (in Japanese; title translation is the author's).