HONG KONG · HKMA

# HKMA's AI Push on *Suspicious Payments*

## Opens a New Governance Gap

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

# HKMA's AI Push on Suspicious Payments *Opens a New Governance Gap*

As AI begins interpreting payment activity, institutions need to govern the meaning behind consequential decisions, not only the data, tools and models involved.

MAUREEN DOYLE-SPARE

Hong Kong is moving AI further into the financial-control environment. On September 16, the government released its first Five-Year Plan for Economic and Social Development 2026-2030 alongside the 2026 Policy Address. In the section on managing new financial risks, it says the Hong Kong Monetary Authority is considering using AI to analyze payment data to strengthen banks' detection of suspicious transactions. The government also plans to consult on a legal framework to improve the detection and removal of fraudulent content, including content produced through misuse of AI. This places AI inside Hong Kong's formal financial-risk agenda and closer to the control processes through which institutions decide what activity may be suspicious, high-risk or reportable.

There is good reason for that move. Transaction monitoring has always required more than data. A payment can be observed directly, but whether it is suspicious, high-risk or reportable depends on the customer, the jurisdiction, the surrounding activity and the institution's own policies. The HKMA's 2024 guidance on the use of AI for monitoring suspicious activities recognized the value of bringing more of that context into monitoring, including customer risk profiles, historical transaction patterns and other relevant information. Its June 2026 work on AI in fighting financial crime pushes the same direction further.

## Suspicion is not a data point

Suspicion is not inherent in a payment. It is resolved from the context around it: the customer, counterparties, transaction history, jurisdiction, relevant typologies and the institution's own risk policies. Richer context can make detection more effective. It also changes what the institution is asking the system to do. Once an AI-enabled system moves beyond finding information and begins resolving what that information means for a particular case, the institution is no longer governing only data quality, model performance or tool permissions. It is governing interpretation. That reasoning-layer concern becomes more relevant in HKMA's cross-sector GenA.I. Sandbox++ environment, where AI use

© 2026 Maureen Doyle-Spare.

cases can span banking, securities, insurance and payments and where conclusions formed in one workflow may affect another. That is the point at which interpretation becomes a runtime governance concern.

This resolved meaning is the Operational Interpretation: the working meaning under which an autonomous system proceeds. Transaction monitoring is a clear example because terms such as suspicious, sanctioned, high-risk and reportable acquire operational meaning only in context. The system has to resolve what those terms mean for the case in front of it.

The exposure is not limited to bad data or a defective model. Systems can each return accurate, policy-conforming results and still contribute to an interpretation the institution never authorized. A transaction-monitoring scenario illustrates the problem: a multi-currency wire is evaluated across an AML platform, sanctions screening and a typology library. Each source can be locally correct while the combined interpretation still misses the institution's authoritative definition of reportable cross-jurisdictional structuring. Nothing needs to break in the conventional sense. The gap can sit in the meaning assembled across otherwise valid systems. That is the Invisible Failure: the surrounding controls can all appear to work while the institution acts under an interpretation it never authorized.

That is why the Hong Kong development matters beyond AML technology. The more effectively AI can combine context, identify relationships and resolve significance, the more useful it becomes to financial-crime teams. The same capability also moves more institutional meaning into the runtime decision. Governance has to keep pace with both sides of that equation.

## The boundary changes when interpretation begins to carry authority

There is a meaningful difference between an AI system that helps an investigator find relevant information and one whose interpretation can influence whether an alert is prioritized, a case is escalated, activity is held or a conclusion is passed into another automated workflow. An AI system can remain within its technical permissions and still proceed under an interpretation the institution never authorized. As its role becomes more consequential, the institution needs to know not only whether the system was allowed to act, but whether the meaning supporting that action remained within authorized bounds.

The runtime architecture separates the interpretation from the point at which it becomes executable. Authorization Emission is the runtime event at which a governed interpretation crosses into execution authority. The distinction matters because semantic validation and execution are not the same event. An interpretation can be evaluated and held before authority is released, or it can be allowed to proceed. In a consequential financial-control workflow, that separation preserves the opportunity to intervene before meaning becomes action, rather than relying only on reconstruction after the fact.

© 2026 Maureen Doyle-Spare.

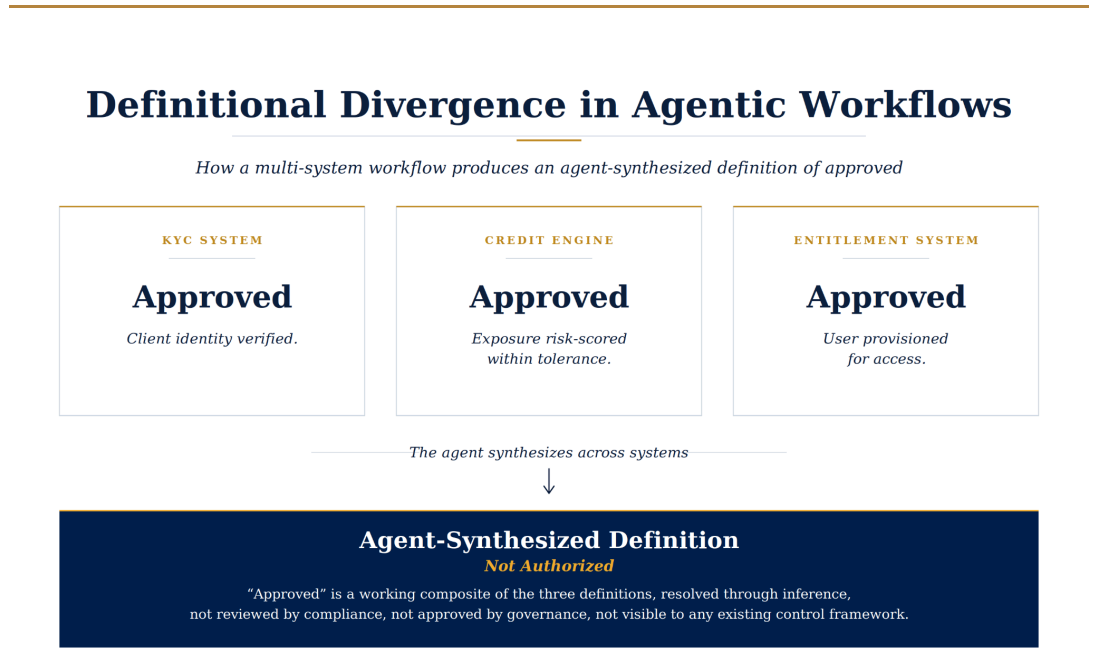



***Figure 1.*** *Definitional Divergence in Agentic Workflows. Source: Doyle-Spare (2026).*

This is also where evidence has to become more precise. A conventional audit trail can show which data were accessed, which model ran and what action followed. It does not necessarily establish that the operational meaning resolved from those inputs matched the institution's authorized meaning at the point the decision became consequential. Within the Semantic Control Plane, Verification of Runtime Semantic Resolution, or V-RSR, evaluates the system's runtime interpretation against the active Reasoning Baseline before the Gate issues its decision. The result can then be preserved as decision evidence. The purpose is not another retrospective log. It is a control point at which the institution can still prevent an unauthorized interpretation from becoming executable authority.

## Not every AI use case needs the same control

This does not require placing a human reviewer in front of every AI-supported decision. That approach would recreate much of the bottleneck the technology is intended to reduce. Financial institutions will use AI across a wide range of AML and financial-crime activities, and those uses do not carry the same exposure. A system that summarizes case information for an investigator is different from one that resolves multiple regulated signals into a conclusion affecting a customer, a filing or another automated process. The Runtime Governance Applicability Test, or RGAT, distinguishes those cases before heavier runtime controls are applied. The closer a system moves to resolving institutional meaning, and the farther that meaning can travel through other processes, the stronger the case for pre-execution governance.

That makes the HKMA's progression particularly interesting. Hong Kong has moved from examining AI-enabled transaction monitoring to supporting broader AI use in suspicious-activity monitoring and financial crime, and now to placing AI analysis of payment data inside its formal policy agenda. None of this means that autonomous financial-crime decisioning is suddenly arriving fully formed. It does mean that the regulatory and operational discussion is moving toward systems that do more

© 2026 Maureen Doyle-Spare.

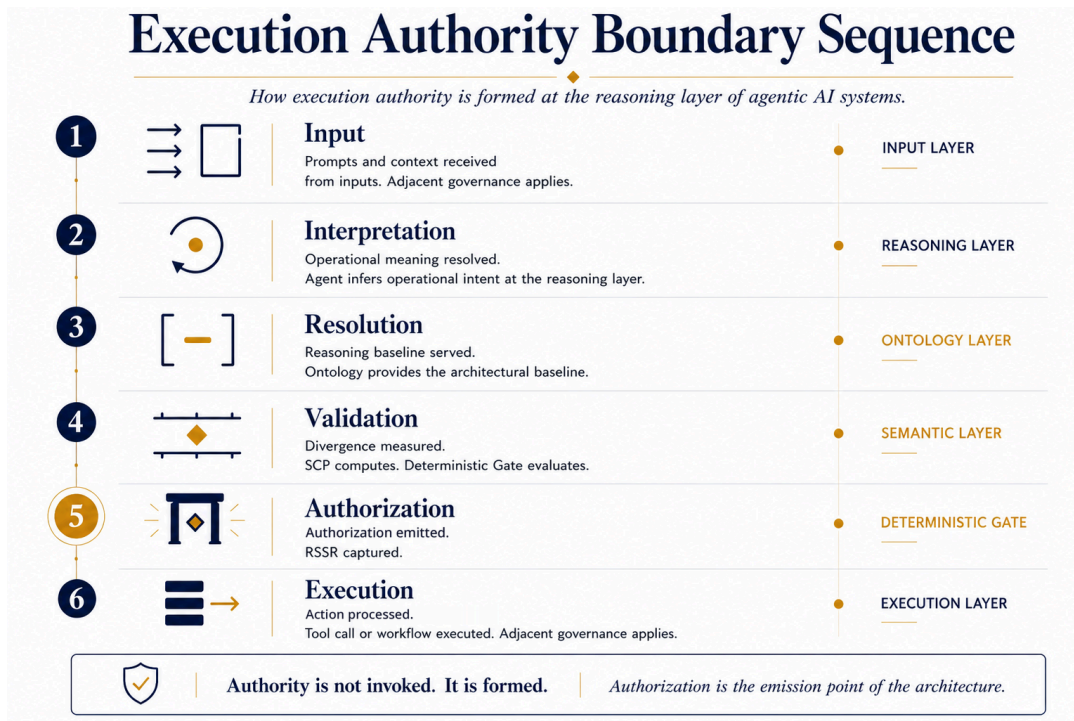



***Figure 2.*** *Execution Authority Boundary Sequence. Reproduced from Doyle-Spare (2026).*

than retrieve, summarize and organize information.

Hong Kong's policy direction is clear. AI is moving closer to the point at which financial institutions interpret activity and decide what it means. That is a positive development when richer context improves detection and helps investigators focus on genuinely higher-risk activity. It also makes the governance problem more exacting, because the institution increasingly has to control not only the information available to the system and the action it is permitted to take, but the operational meaning that connects the two.

For financial institutions, this is the next control boundary. As AI moves deeper into suspicious-activity monitoring, the institution needs to establish, before a consequential decision proceeds, that the Operational Interpretation behind it remains within authorized meaning. Hong Kong's new policy direction does not yet formalize that control. It makes the need for it more practical. The governance issue is no longer only whether AI can identify suspicious activity. It is whether the institution can govern the meaning AI resolves before that meaning carries financial-control authority. ◆

## References


1 *Hong Kong SAR Government. The Chief Executive's 2026 Policy Address (Management of New Financial Risks). September 16, 2026. policyaddress.gov.hk/2026/en/policy.html*

2 *Hong Kong Monetary Authority. "Use of Artificial Intelligence for Monitoring of Suspicious Activities." September 9, 2024. brdr.hkma.gov.hk/eng/doc-ldg/docId/getPdf/20241122-3-EN/20241122-3-EN.pdf*

3 *Hong Kong Monetary Authority. "Supporting Adoption of Artificial Intelligence in Fighting Financial Crime." June 22, 2026. brdr.hkma.gov.hk/eng/doc-ldg/docId/20260622-1-EN*

4 *Doyle-Spare, Maureen. "Hong Kong HKMA: Reasoning Layer Governance — Public Comment, April 2026." The Agentic AI Governance Playbook, Substack. maureendoylespare.substack.com/p/hong-kong-hkma-reasoning-layer-governance-2026-04-18*

5 *Doyle-Spare, Maureen. Agentic AI Systems Governance: A Runtime Reference Architecture for the Reasoning Layer and the Semantic Control Plane in Regulated Financial Institutions. Zenodo, 2026. doi.org/10.5281/zenodo.20749051*


© 2026 Maureen Doyle-Spare.

6 *Doyle-Spare Research Library.* *maureendoylespare.com/research*

---

Originally published in *The Agentic AI Governance Playbook* on Substack, September 18, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.




*The governance issue is no longer only whether AI can identify suspicious activity. It is whether the institution can govern the meaning AI resolves.*

THE AGENTIC AI GOVERNANCE PLAYBOOK

## Download the research

Papers, briefs and reference architecture at maureendoylespare.com

Visit the Playbook

WEEKLY NEWSLETTER

## Follow the newsletter

The Agentic AI Governance Playbook Weekly Newsletter

Follow Newsletter

**INTELLECTUAL PROPERTY:** © 2026 Maureen Doyle-Spare. All rights reserved. No part of this article, including its text, figures, tables, and images, may be reproduced or distributed without the author's prior written permission.