

AI agents in trade execution, market surveillance and tokenised infrastructure debt

Three uses of AI that arrived before their controls.
Terence Kok with Dr. Christina Chua.

Each use raises one control question

1

Agents that place orders

Who is liable when an agent trades outside its mandate?

2

LLM surveillance

How do you prove a quieter alert queue is a better one?

3

Tokenised infrastructure debt

Can live asset data be trusted to value the bond?

\$460m

**lost in about 45 minutes when old,
unused trading code went live on one
server**

More than 4 million executions in 154 stocks. The SEC fined the firm \$12 million. Source: SEC order 34-70694, 2013.

The firm running the agent is the adviser



Financial Advisers Act, s36

A recommendation needs a reasonable basis. Loss from reliance on one that fails the test gives a right to damages.



Third-party agents

When a client connects their own agent, how loss is shared is still unsettled.



MAS AI guidelines

Proposed November 2025, still in draft, with a 12-month transition after issue.

Four graded stops outside the agent

1

Pause the agent

Stops proposing orders. Positions stay.

2

Block new orders

The gateway rejects the agent's credentials.

3

Cancel and revoke

All open orders cancelled, access revoked.

4

Firm-wide kill

All automated flow stops at market access.

2 yrs+

that some firms ran faulty surveillance alerts without knowing

FCA Market Watch 79, May 2024. A shorter alert queue looks like efficiency. It is also what a system that has stopped catching abuse would produce.

Measure what the model misses

- 1** Freeze the old numbers
- 2** Run both systems in parallel
- 3** Seed known cases
The share caught estimates recall.
- 4** Sample below the line
Zero finds in 3,000 items: miss rate under 1 in 1,000.
- 5** Adjudicate the disagreements

Use asset data to **trigger revaluation**

Operator data has an interested owner, a live model is still a Level 3 model, and continuous prices on collateral create margin calls. Publish the data, agree thresholds, and keep an independent valuer on the price.

Five questions to ask

- 1** Which AI systems can place, change or cancel orders, and who owns each?
- 2** How long to stop an agent at 3 a.m., and when was that tested?
- 3** Can our logs show why an agent placed a given order?
- 4** What are our seeded-case catch rate and below-the-line miss rate?
- 5** Who produces and who checks the data behind a tokenised asset's value?



**The stop must act faster
than the thing it is meant
to stop.**



Terence Kok · AI & Innovation Strategist

SAVE · COMMENT · FOLLOW