[redacted]

From: Sarthak Gupta [redacted]
Sent: Wednesday, 19 August 2026 4:21 PM
To: Policy Development
Subject: Submission on the Governance Review, draft CPS 510, Question 6 (AI governance)

EXTERNAL EMAIL: Do not click any links or open any attachments unless you trust the sender and know the content is safe.

Attention: General Manager, Policy and Frameworks, Policy and Advice Division, Australian Prudential Regulation Authority

Dear General Manager,

I am writing in response to the Governance Review consultation and the draft Prudential Standard CPS 510 Governance, specifically on Question 6 on artificial intelligence governance. I work in model risk and quantitative finance, where the supervision of models used in prudential decisions has been under formal expectations for more than a decade, and I want to offer one structural point drawn from that experience.

The draft treats AI oversight well, but as its own subject. My suggestion is the opposite: where an AI or agentic system materially informs a prudential decision, the board's governance of it should inherit the accountability, independent-review and reporting obligations that already apply to models, rather than sit in a parallel AI track. A separate track sounds tidier, but it creates a gap at the seam, where a system that would be governed as a model if a statistician built it escapes that governance because it is labelled AI. The cleaner requirement is coverage by function, not by technology name.

Two expectations would give that principle teeth.

First, independence of review should scale with the materiality of the decision the system informs, not be a single checkbox. A model that sizes a capital position and a model that routes a low-value query should not attract the same intensity of challenge, and the standard is stronger if it says so.

Second, an agentic system that can take an action, not only produce a score, needs a named accountable owner and a written authority boundary: what it may decide on its own, and what must return to a person. That boundary, and the traceability of the decisions it takes inside it, is what a board can actually oversee. Without it, "human oversight" has nothing specific to attach to.

I would be glad to expand on any of this if it is useful. I am content for this submission to be published on the APRA website.

Best regards,
Sarthak Gupta
[redacted]