

The Assumption Every AI Governance Framework Makes — and Why It Fails in GCC

How informal authority, board composition constraints, and regulatory fragmentation undermine AI governance frameworks in GCC institutions

Majid Mumtaz · CIA · ACA · FCCA | Director of Internal Audit | 20+ Years GCC

EXECUTIVE SUMMARY

Every AI governance framework published in the last three years — NIST AI RMF, ISO 42001, the UAE AI Policy Framework — is built on a shared assumption: that formal authority structures are real. That the person named “AI Oversight Committee Chair” has actual decision authority. That the board which approved the AI governance policy can evaluate whether it works. That regulatory jurisdiction maps cleanly onto where AI systems operate.

In mature governance markets, this assumption is close enough to true that frameworks built on it are functional. In GCC, it is not close enough. And AI does not expose this slowly.

This paper identifies three structural realities that generic frameworks were not designed to handle — and specifies the accountability reallocation each requires.

WHAT GENERIC FRAMEWORKS ASSUME AWAY

NIST AI RMF, ISO 42001, and their derivatives share a structural logic: define roles, assign responsibilities, test fulfilment, escalate failures up the governance hierarchy. This logic works when formal roles carry real authority, when the governance hierarchy is the actual decision chain, and when the people at the top have the capacity to evaluate what is escalated to them.

The framework does not account for what happens when none of these conditions hold simultaneously. In GCC family-controlled entities, all three fail at once.

STRUCTURAL REALITY 1 — INFORMAL AUTHORITY CONCENTRATION

In a GCC family-controlled entity, the founder or senior family principal makes consequential decisions. What is GCC-specific is the combination: decision authority that is informal (not mapped to any role), concentrated (one or two individuals), and structurally immune to governance architecture — no board resolution, no documented rationale, no audit trail.

Governance frameworks are designed for entities where authority flows from documented delegation. In GCC family businesses, authority flows from relationship. The CAE who has learned this quickly: the organisation chart describes hierarchy, not who decides.

What AI amplifies:

AI operates at speed. The informal veto that previously slowed consequential decisions — no meeting scheduled, pending review, timing never right — loses its friction when an AI system generates a recommendation and expects a response. The informal authority is still there. It is now being asked to intervene faster, on outputs it did not commission and may not understand.

From Practice: At a GCC retail holding group, an AI-driven cash flow forecasting model flagged three subsidiary businesses as structurally loss-making and recommended divestment. The model's output was reviewed by the formally appointed CFO and documented as "considered." The businesses were retained — the decision made by the family principal informally, with no governance trace. Eighteen months later, two of the three required emergency liquidity support. The governance record showed the AI framework had been followed throughout. No finding was raised because the formal accountability owner had formally considered the recommendation. The person who made the actual decision was not in the accountability structure.

Where informal authority is total — where the family principal both owns the consequence and holds the decision — no governance framework changes the accountability structure. The correct finding is not a control gap. It is a governance design limitation the board needs to understand exists.

What actually needs to change:

Map actual decision authority — not org chart authority. In every consequential AI use case, identify who will informally ratify or veto the output. Assign consequence at that level. Ask: who has the ability to override this system's output, and are they accountable for that override? If the answer is "the founder, informally, without documentation" — that is a governance design constraint to design around, or to surface as an unresolvable finding.

STRUCTURAL REALITY 2 — BOARD COMPOSITION AND THE ANALYTICAL CAPACITY GAP

GCC boards, particularly in family-controlled entities and pre-IPO companies, frequently include members whose appointment reflects relationship, not competence. Board composition in concentrated ownership structures serves multiple functions simultaneously — governance, relationship management, stakeholder signalling. The governance function is often not the primary one.

Governance codes require boards to ensure AI risk is identified, assessed, and managed. This requires someone who can interrogate model assumptions, evaluate training data quality, assess output interpretation, and challenge risk appetite statements about AI-specific exposures. In most GCC boards, this person does not exist — not because boards are unintelligent, but because the composition process was not designed to generate this capability.

What AI amplifies:

AI creates opacity. The more capable the system, the harder it is to explain what it is doing in terms a board can evaluate. Governance frameworks respond by requiring "explainability" — a technical property that does not, by itself, create board analytical capacity. The gap between what boards are formally accountable for and what they can actually evaluate widens with every deployment.

When an AI-assisted decision fails, the board is accountable for a governance failure it could not have detected. Formal accountability and actual decision authority are misaligned — invisible until something breaks.

What actually needs to change:

Stop treating board AI literacy training as the solution. Training does not close the analytical capacity gap in time for the AI systems already in production. Redefine board AI oversight not as understanding, but as challenge capacity: who is the independent technical voice in the room that interrogates AI outputs on the board's behalf? Not the Chief AI Officer — that is a management function with a structural conflict of interest. An independent technical advisor to the audit committee, with a clear mandate, a specific challenge protocol, and accountability that makes oversight real rather than formal.

In the absence of this, board AI governance oversight is a documentation exercise.

STRUCTURAL REALITY 3 — REGULATORY FRAGMENTATION AND THE COMPLIANCE SURFACE PROBLEM

A GCC entity operating in both UAE and KSA faces two distinct regulatory environments for AI. The UAE operates under the TDRA AI Policy Framework, with DIFC and ADGM providing separate guidance for regulated entities. In KSA, the Saudi Data and AI Authority (SDAIA) has published National AI Ethics Principles, and the PDPL establishes data handling requirements for AI training data. These frameworks share language — accountability, transparency, fairness — without sharing definitions, enforcement mechanisms, or jurisdictional boundaries.

What AI amplifies:

AI systems do not respect regulatory boundaries. A model trained on UAE customer data, deployed to KSA operations, making recommendations about KSA customers creates a compliance surface neither jurisdiction has fully addressed. The entity is simultaneously subject to both frameworks and definitively covered by neither. Regulatory clarity typically follows technology deployment by 18 to 36 months. AI deployment cycles in GCC are compressing.

The governance framework response — a jurisdiction-by-jurisdiction compliance matrix — creates a compliance artefact, not compliance. The matrix tracks regulatory text, not where AI systems are operating or what data they are consuming across borders in real time.

What actually needs to change:

Govern to the highest-restriction jurisdiction. Single standard, applied consistently, with a named accountability owner responsible for maintaining it as regulations evolve. Less sophisticated than a multi-jurisdiction matrix; more likely to produce actual compliance.

Audit the AI system, not the AI policy. Map where every system in production is making decisions, what data it is consuming, what jurisdictions that data originates from, and who is accountable for each decision it makes. This is the precondition for knowing whether your regulatory compliance claims are true.

THE DEPLOYMENT PRESSURE PROBLEM

Both UAE and KSA national AI strategies create organisational pressure to demonstrate adoption. The pressure to deploy is institutional and externally signalled. The pressure to govern well is internal and self-generated. When these conflict — as they increasingly do in pre-IPO GCC entities — deployment pressure carries more organisational momentum.

The CAE's role: establish what "AI governance quality" means internally before it is externally defined. The entity that waits for the regulator to define it will be governing reactively, with systems already in production not designed with the required accountability structure. The entity that defines it first has a defensible position and the time to build it properly.

WHAT BOARDS AND CAES SHOULD DO DIFFERENTLY

For Boards:

Commission an AI accountability map — not a policy, a map. For every AI system in production: who makes decisions based on its outputs? Who is accountable if those outputs are wrong? If the answer is "the system" or "the team," the accountability structure is incomplete.

Define independent technical challenge capacity. A named function, a specific mandate, a documented challenge protocol. Not a training programme.

Approve a single standard for cross-jurisdictional AI governance. Accept the compliance cost of the most restrictive jurisdiction. The alternative cost is higher.

For CAEs:

Audit the decision chain, not the governance framework. Test whether the named accountability owners actually received, evaluated, and acted on AI outputs in the last reporting period. Document the gap between formal and actual decision authority — and where that gap is structural and unresolvable, say so explicitly.

Treat informal authority as a governance design input, not a control deficiency. If the real decision-maker is not in the governance framework, the framework does not govern the decision. Escalate this as a structural finding to the audit committee.

Establish AI audit scope before deployment scales further. The AI audit universe is expanding faster than audit capacity in most GCC functions. Define scope now, or scope will be defined by the next governance failure.

CLOSING OBSERVATION

The organisations that will handle AI governance well in GCC are not the ones with the most sophisticated frameworks. They are the ones that understand where their formal governance structures actually stop working — and design accountability for that specific gap rather than adding documentation to avoid naming it.

AI governance frameworks do not fail because they are wrong. They fail because they assume a governance maturity that GCC institutions, in many cases, have not yet built. And they fail silently: the compliance record is clean, the policies are approved, the committee has met — and the actual decisions are happening somewhere the framework was never designed to reach.

Building that maturity is not a technology question. It is an accountability question. That has always been the harder one.

About the Author: Majid Mumtaz is a Director of Internal Audit with 20 years of GCC experience across listed, pre-IPO, and family-controlled entities in UAE and KSA. Credentials: CIA (IIA), ACA (ICAEW), FCCA (ACCA).

Contact: majidrajpar@gmail.com | [linkedin.com/in/majid-m-4b097118/](https://www.linkedin.com/in/majid-m-4b097118/)