MAHMOOD.INSIGHT

THE FUTURE OF FINANCIAL INFRASTRUCTURE

AI in the Omani Banking System

The technology that is now both the strongest defence against financial crime and its fastest-growing source — and the control gap every bank must close.

Edition 02
Q4 2026

Coverage
Sultanate of Oman

Series
Quarterly

Mahmood Al-Hattali · Independent Strategic Advisor · Banking & Payments

IN THIS EDITION

Contents

01 The headline number — AI's double edge 04

02 The framework lens — the AI Control Gap 05

03 Supervision — how regulators are splitting 06

04 The contrarian read — the risk of going slow 07

05 Three predictions for next quarter 08

EXECUTIVE SUMMARY

Five things every CEO should know this quarter.

1 **AI is now both shield and sword.** Generative-AI-enabled fraud in the U.S. is projected to rise from $12.3bn (2023) toward $40bn by 2027 — while AI fraud defence is cutting false positives by up to 60% at major banks.

2 **Adoption is no longer the question.** Organisational AI adoption has reached 88% globally, and generative AI is used in at least one business function at 70% of organisations. The laggard's risk is now larger than the early mover's.

3 **The danger is the control gap.** When AI adoption outpaces AI governance, the space between them is exactly where model risk, bias, and undetected fraud live. This edition names and maps that gap.

4 **Regulators are split.** Some supervisors are tightening scrutiny of AI in credit and fraud; others warn the real risk is inaction — leaving banks under-armed against criminals who already use AI. Oman's posture is still being written.

5 **Governance is the competitive asset.** In a market that doubled its payment volume (Edition 01), the banks that pair fast AI adoption with disciplined control will both lose less to fraud and earn supervisory trust faster.

FROM THE DESK

The same tool now commits the crime and catches it.

Edition 01 showed a market that crossed into digital faster than its institutions were built for. This edition is about what is now arriving on those same rails: artificial intelligence, on both sides of the ledger.

There is a symmetry to AI in banking that makes it unlike previous technology waves. The model that detects a fraudulent transfer in milliseconds is a close cousin of the model that generates the deepfake voice authorising it. Defence and attack are drawn from the same well. That is why this cannot be read as a simple "adopt AI and win" story — it is an arms race, and the side that governs its weapons best wins.

My background is forensic. I spent years distinguishing a red flag from a conclusion in financial reporting. That discipline is exactly what AI in banking now demands at scale: the technology can surface ten thousand flags a second, but judgement about which are real, and accountability for the ones missed, cannot be automated away. This edition reads AI through three lenses — fraud and financial crime first, then credit and model risk, then supervision itself.

> THE THROUGHLINE
>
> *The winning Omani bank will not be the one that adopts AI fastest, nor the one that governs it most cautiously — it will be the one that closes the gap between the two.*

A note on sourcing

Global AI figures derive from the Stanford HAI AI Index 2026 and PwC's 2026 Global AI Jobs Barometer; supervisory context from the BIS, ECB, and OCC; fraud figures from named industry and institutional sources. The methodology page tiers each. Proprietary frameworks are labelled as the author's analytical models, and predictions are graded on a defined scale (see page 8).

01 · THE HEADLINE NUMBER

AI's double edge, in two numbers.

The same technology is the fastest-growing threat in financial crime and the most powerful defence against it. Both are true at once.

THE SWORD

$40bn

Projected U.S. generative-AI-enabled fraud losses by 2027, up from $12.3bn in 2023 — a ~32% compound annual growth rate.[1]

THE SHIELD

−60%

Reduction in false positives reported by a major global bank after deploying AI-driven risk assessment; others report up to 300% better detection rates.[2]

The defensive case is no longer speculative. Roughly 90% of financial institutions now use AI for fraud detection, and major institutions report material results — fewer false positives, faster detection, lower investigation cost. On the threat side, generative AI has lowered the barrier to entry for criminals: synthetic identities, deepfaked authorisations, and social-engineering at scale that rule-based systems were never built to catch.

Why this matters more in Oman than elsewhere

Edition 01 documented a market whose digital payment volume more than doubled in a year. Fraud surface scales with transaction volume — so a market growing this fast is, by definition, expanding its attack surface faster than most. The strategic read: in a high-growth digital market, under-investing in AI defence is not caution. It is exposure.

> THE QUESTION FOR EVERY OMANI BANK
>
> *The question is no longer "should we use AI against fraud." Criminals already are. It is "is our AI defence improving at least as fast as the AI attacking us?"*

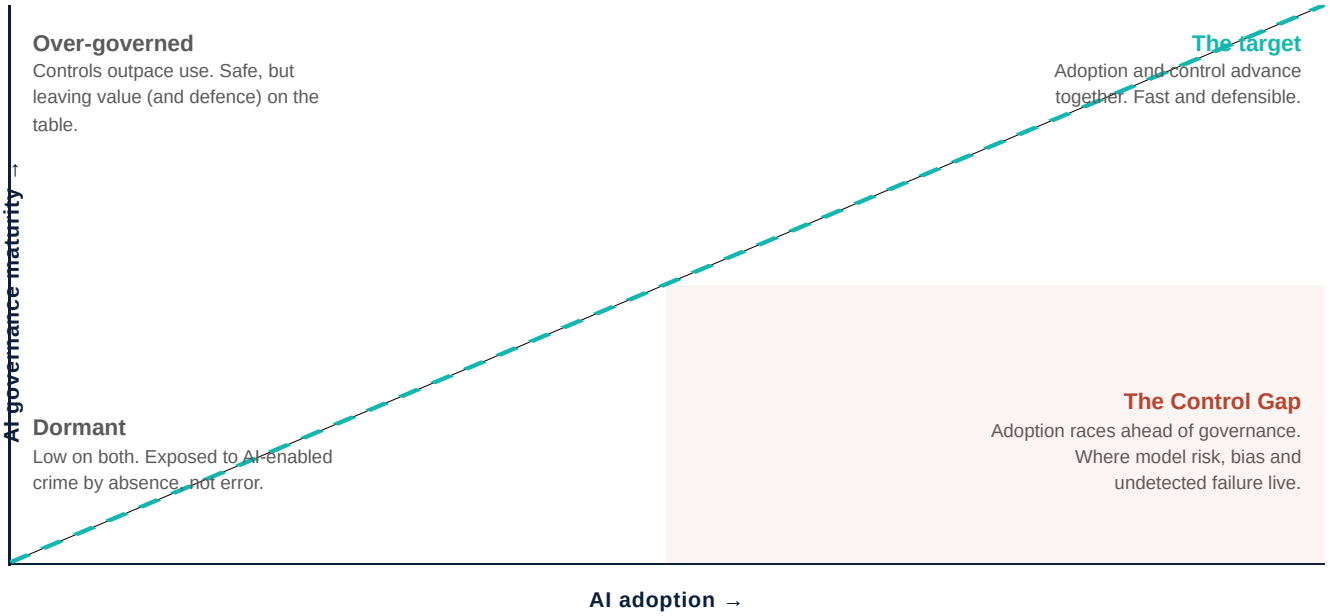
1. Deloitte Center for Financial Services projection, as reported across industry analyses (2025–26); baseline $12.3bn (2023).

2. HSBC Dynamic Risk Assessment (−60% false positives) and Mastercard generative-AI detection gains; company and industry reporting, 2025–26.

02 · THE FRAMEWORK LENS

The AI Control Gap.

A diagnostic, not a description. Plot any bank on two axes — how much AI it has adopted, and how mature its governance of that AI is. The distance from the diagonal is the risk.



Most banks, under competitive and fraud pressure, move **right before they move up** — they adopt AI in fraud, credit, and service faster than they build the governance to control it. That trajectory drops them into the bottom-right: the Control Gap. It is the most dangerous quadrant precisely because it feels like progress.

Reading it with the Four Layers

This connects directly to the Four Layers of Financial Infrastructure from Edition 01. Adoption is mostly a *Rails* problem — buyable. Governance is a *Data, Org and Culture* problem — which is why it lags. Banks close the gap not by buying more AI, but by maturing the three layers that cannot be purchased.

> **THE DECISION THE FRAMEWORK FORCES**
>
> Don't ask "how much AI have we adopted?" Ask "how far right of our governance line are we — and what is the single control we'd add this quarter to move up, not just right?"

03 · SUPERVISION

How the world's regulators are splitting.

There is no settled supervisory consensus on AI in banking yet. The major authorities have visibly diverged — and where Oman's CBO lands is still an open, consequential question.

- **The tightening camp.** The ECB stepped up monitoring of AI in credit scoring and fraud detection through 2025, via dedicated data collections and targeted supervisory dialogue, and folded AI into its 2026–28 supervisory priorities on operational resilience.[3]
- **The integrating camp.** Germany's BaFin reframed AI not as a separate "ethics" question but as part of ICT risk management — AI must be understood, controlled, and audited like any other critical system.[4]
- **The accelerating camp.** In the U.S., the Fed, OCC and FDIC narrowed model-risk guidance so it no longer constrains generative and agentic AI — explicitly to avoid slowing adoption.[5]
- **The supranational anchor.** The BIS published direct guidance on the governance of AI adoption in central banks, giving regulators a common reference for SupTech, AML/CFT analytics, and payment-system oversight.[6]

What this means for the CBO

A regulator entering AI supervision today is not a laggard — it is arriving in step with a global community that is itself still deciding. That is an advantage: Oman can adopt the BIS reference frame and the clearest elements of each camp without inheriting any one jurisdiction's mistakes. The forensic point: explainability and auditability are the common thread across every camp. A bank that cannot reconstruct why its model declined a loan or flagged a transfer is exposed under all of them.

> THE SUPERVISORY READ
>
> *Whichever camp the CBO joins, the demand will be the same: show your work. Banks that build auditable AI now will clear whatever bar arrives — those that don't will retrofit it under pressure.*

3. ECB Banking Supervision, AI in credit scoring and fraud detection; Supervisory Priorities 2026–28 (2025–26).

4. BaFin guidance on ICT risks associated with AI use, December 2025.

5. Federal Reserve / OCC / FDIC revised Model Risk Management guidance (SR 26-2), April 2026.

6. Bank for International Settlements, "Governance of AI adoption in Central Banks," 2025.

04 · THE CONTRARIAN READ

The bigger risk may be going too slow.

The instinctive supervisory posture toward AI is caution. This section argues the opposite case deserves a hearing — because the threat is not waiting for the governance debate to conclude.

The conventional view

Slow down. AI is opaque, models can be biased, generative systems hallucinate. Require months of model-risk review before deployment, demand full explainability, treat each use case as high-risk until proven otherwise. Caution protects the customer and the system.

The case against it

While banks deliberate, criminals deploy. Generative-AI fraud is compounding at roughly a third per year, and the attackers face no model-risk committee. A bank that takes a year to approve a defensive AI tool has given the adversary a year of unanswered advantage. In that framing, excessive caution is not safety — it is a different, slower-moving risk.

THE ASYMMETRY

The attacker ships AI instantly, iterates daily, faces no governance.

The defender reviews for months, documents everything, deploys cautiously.

A defence that updates slower than the attack is not a defence. The governance challenge is to be both fast and accountable — not to choose.

WHAT MOST BANKS WILL GET WRONG

They will treat AI governance as a brake — a committee that says no slowly. The bank that wins treats governance as a clearance lane: a fast, disciplined process that lets defensive AI ship in weeks with full auditability, instead of dying in review while fraud compounds. Speed and control are not opposites here. Failing to build both is the real exposure.

05 · PREDICTIONS FOR NEXT QUARTER

Three calls — specific enough to be graded.

Falsifiable on purpose. They feed the annual scorecard, graded on the same defined scale as Edition 01.

PREDICTION 01

An Omani or GCC bank publicly announces an AI-based fraud or AML capability.

With fraud surface scaling alongside payment volume and global peers reporting strong results, a visible AI-fraud or AML deployment is the most likely next headline in the region's banking AI story.

PREDICTION 02

A GCC regulator issues or signals AI / model-risk guidance for banks.

As the BIS reference frame spreads and supervisors elsewhere move, at least one Gulf authority is likely to publish, consult on, or publicly flag forthcoming AI governance expectations.

PREDICTION 03

Generative-AI fraud (deepfake or synthetic-identity) is publicly reported in the region.

Given the documented global surge and the region's high digital adoption, a publicly disclosed generative-AI fraud incident or attempt in the GCC becomes increasingly likely.

HOW THESE WILL BE GRADED

*Each prediction will be scored in the annual outlook on a defined scale — **Confirmed**, **Partially confirmed**, or **Not confirmed** — publicly, including any misses. Stating the scale in advance is what makes the scorecard a method rather than a marketing exercise.*

This is Edition 02.

The State of Omani Banking & Payments is published quarterly, reading each quarter's developments through the same lens: where the market is going, where it ends up, and what institutions must build to get there.

Next edition: Q1 2027 — the year-one prediction scorecard.

SOURCES & METHODOLOGY

Every figure, traced.

Each data point carries its source, year, and a tier: P = primary institutional/official publication, S = secondary or industry reporting.

Claim	Year	Source	Tier
Gen-AI fraud: $12.3bn (2023) → ~$40bn (2027 proj.)	2023–27	Deloitte projection, via industry reporting	**S**
−60% false positives (AI risk assessment)	2025	HSBC; company reporting	**S**
Up to 300% detection improvement (gen-AI)	2025	Mastercard; company reporting	**S**
~90% of institutions use AI for fraud detection	2025	Feedzai AI Trends Report	**S**
Organisational AI adoption 88%; 70% gen-AI in a business function	2025	Stanford HAI AI Index 2026	**P**
AI wage premium ~62%; two-track labour market	2025	PwC 2026 Global AI Jobs Barometer	**P**
ECB stepped-up AI monitoring; 2026–28 priorities	2025–26	ECB Banking Supervision	**P**
BaFin: AI as ICT risk	Dec 2025	BaFin guidance	**P**
Model-risk guidance narrowed for gen/agentic AI	Apr 2026	Fed/OCC/FDIC (SR 26-2)	**P**
Central-bank AI governance reference	2025	Bank for International Settlements	**P**

Method

- **Primary where it counts.** The macro AI, workforce, and supervisory claims are cited to primary institutional publications (P). Fraud performance figures, by nature proprietary to banks and vendors, are industry-reported (S) and framed as illustrative, not audited.
- **Global evidence, local lens.** International data is used to frame the Omani question; it is not presented as Omani data. Region-specific claims are flagged as such.
- **Frameworks are interpretive.** The AI Control Gap and the Four Layers of Financial Infrastructure are the author's analytical models, not data.
- **Predictions are falsifiable by design** and graded Confirmed / Partially confirmed / Not confirmed in the annual outlook, including any misses.

This report synthesises publicly available information for analytical and discussion purposes. It is not investment, legal, or regulatory advice. Figures are accurate to the author's knowledge at time of publication; readers should consult primary sources for definitive data. © Mahmood Al-Hattali, 2026.

DEFINITIONS

The terms, in one line each.

So the report reads cleanly for specialists and non-specialists alike. Each term as used in this publication.

Term	Definition
Generative AI	AI that creates new content — text, voice, images — rather than only classifying existing data; the basis of both modern productivity tools and deepfake fraud.
Agentic AI	AI that takes actions autonomously toward a goal, not just answering prompts; raises governance questions because it acts before a human reviews each step.
Model risk	The risk that a model is wrong, misused, or misunderstood — and that decisions made on its output cause loss or harm.
Explainability	The ability to reconstruct and articulate why a model produced a given output; central to supervisory acceptance of AI in credit and fraud.
False positive	A legitimate transaction wrongly flagged as fraud; high rates frustrate customers and waste analyst time — a key cost AI aims to cut.
Deepfake	Synthetic audio or video generated by AI to impersonate a real person; increasingly used to authorise fraudulent transfers.
Synthetic identity	A fabricated identity combining real and fake data to pass verification; a fast-growing, AI-accelerated fraud vector.
SupTech	Supervisory technology — AI and data tools used by regulators themselves to monitor banks and payment systems.
AML / CFT	Anti-money-laundering / combating the financing of terrorism — the controls AI is increasingly used to automate and strengthen.

A NOTE ON USAGE

Where this report uses a proprietary framework — the AI Control Gap, the Four Layers of Financial Infrastructure — it is named as such. These are analytical lenses developed for this publication, distinct from the industry terms defined above.