


EDITION 01

From Insight to Foresight

Corporate Governance in the Age of Predictive AI

This report reframes predictive AI not as a technology trend, but as an immediate corporate governance responsibility across sectors.



CORPORATE GOVERNANCE IN THE AGE OF PREDICTIVE AI

Since 2010, the Pearl Initiative has championed good corporate governance as a practical tool for business leadership in the Gulf region. Through action-oriented programmes, applied research, and youth engagement, we work with corporate leaders and cross-sector partners to embed the highest standards of integrity, accountability, diversity, and transparency across the region.

Corporate governance is not static. Each era brings new pressures to the boardroom. The challenges left unexamined today often become the governance failures of tomorrow.

In this context, we are pleased to launch the **Corporate Governance Insights Series**: a platform to examine the emerging issues shaping corporate governance across the region, before they crystallise into systemic risks. Each edition brings together perspectives from across sectors and organisational types, grounded in both practical experience and regional context.

Our first edition focuses on predictive AI, one of the most consequential shifts now entering the boardroom, and one of the least governed. Across the GCC, organisations are embedding AI-driven systems into core decision-making at a pace that is outstripping the governance structures meant to hold those decisions to account. This edition is designed to help boards ask the right questions, understand the risks, and build the foundations needed to govern what is, for many, still largely uncharted territory.

How to Read This Report

This report brings together perspectives across sectors and organisational types, from large corporates and government systems to MSMEs and philanthropic institutions. While the applications of predictive AI differ, the governance challenge remains consistent: ensuring accountability for decisions influenced by systems that are not always fully visible or understood.

Each section explores this challenge through a different lens, highlighting both shared risks and context-specific considerations. Together, they provide a practical view of what governing predictive AI looks like in practice.

CONTRIBUTORS



EDITOR'S NOTE

Corporate governance is not a design exercise that can be set once and left on autopilot. It requires continuous vigilance, active oversight, and the willingness to adapt as conditions change.

That need for evolution has never been more urgent. Business is being reshaped by rising expectations around accountability, transparency, and long-term responsibility. What was once considered good practice is now being actively tested by stakeholders, regulators, and the public. In many cases, the pace of change is now outpacing the governance structures intended to oversee it.

Governance is no longer an internal exercise. It is increasingly visible, tested, and judged in real time.

This is the context in which this series was developed. It is designed as a platform to examine the issues reshaping governance across sectors and organisational types. It also creates space to reflect on what is working and where approaches are falling short, recognising that both are essential to strengthening governance over time.

Each edition focuses on a specific shift within the business ecosystem. This first report looks at predictive AI, one of the most immediate and consequential changes entering the boardroom today. Its influence on decision-making, scale, and speed is forcing organisations to reassess not just what they do, but how those decisions are governed.

At the Pearl Initiative, we see this series as an opportunity to bring together practice-based perspectives from across the region and contribute them meaningfully to the global conversation. The GCC is not only adopting these changes at pace but is also uniquely positioned to shape how governance evolves alongside them.

The challenge now is to ensure governance evolves at the same pace as the systems it is meant to oversee.



Noor Zuberi

Senior Associate -
Governance in Philanthropy
Programme

Pearl Initiative

EDITORIAL SECTION

WHY PREDICTIVE AI IS A CORPORATE GOVERNANCE ISSUE RIGHT NOW

The AI revolution at the consumer level can broadly be understood across two main categories: generative AI and predictive AI.

1. Generative AI

Generative AI refers to systems that create new content in response to user prompts. These models can produce text, images, audio, video, or code by identifying and applying patterns learned from large training datasets. Rather than retrieving existing material, they generate original outputs based on underlying statistical relationships (IBM, 2024).

2. Predictive AI

Predictive AI uses statistical analysis and machine learning to identify patterns in historical data and forecast future outcomes. By analysing past trends and behaviours, these models estimate the likelihood of specific events or risks, supporting decision-making across areas such as customer behaviour, operational risk, demand forecasting, and financial performance (IBM, 2024).

Predictive AI is often used alongside related analytical approaches:

Predictive analytics: estimating what is likely to occur based on historical data

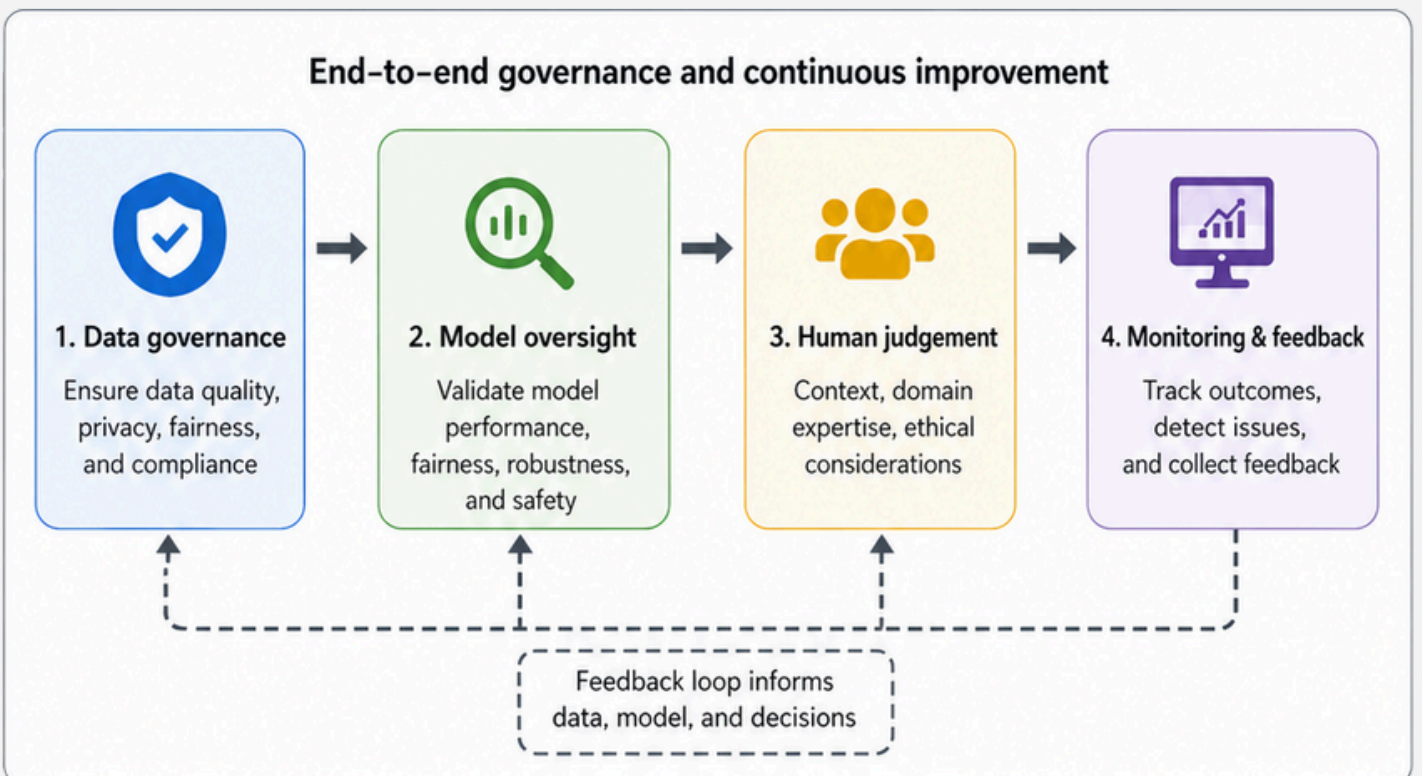
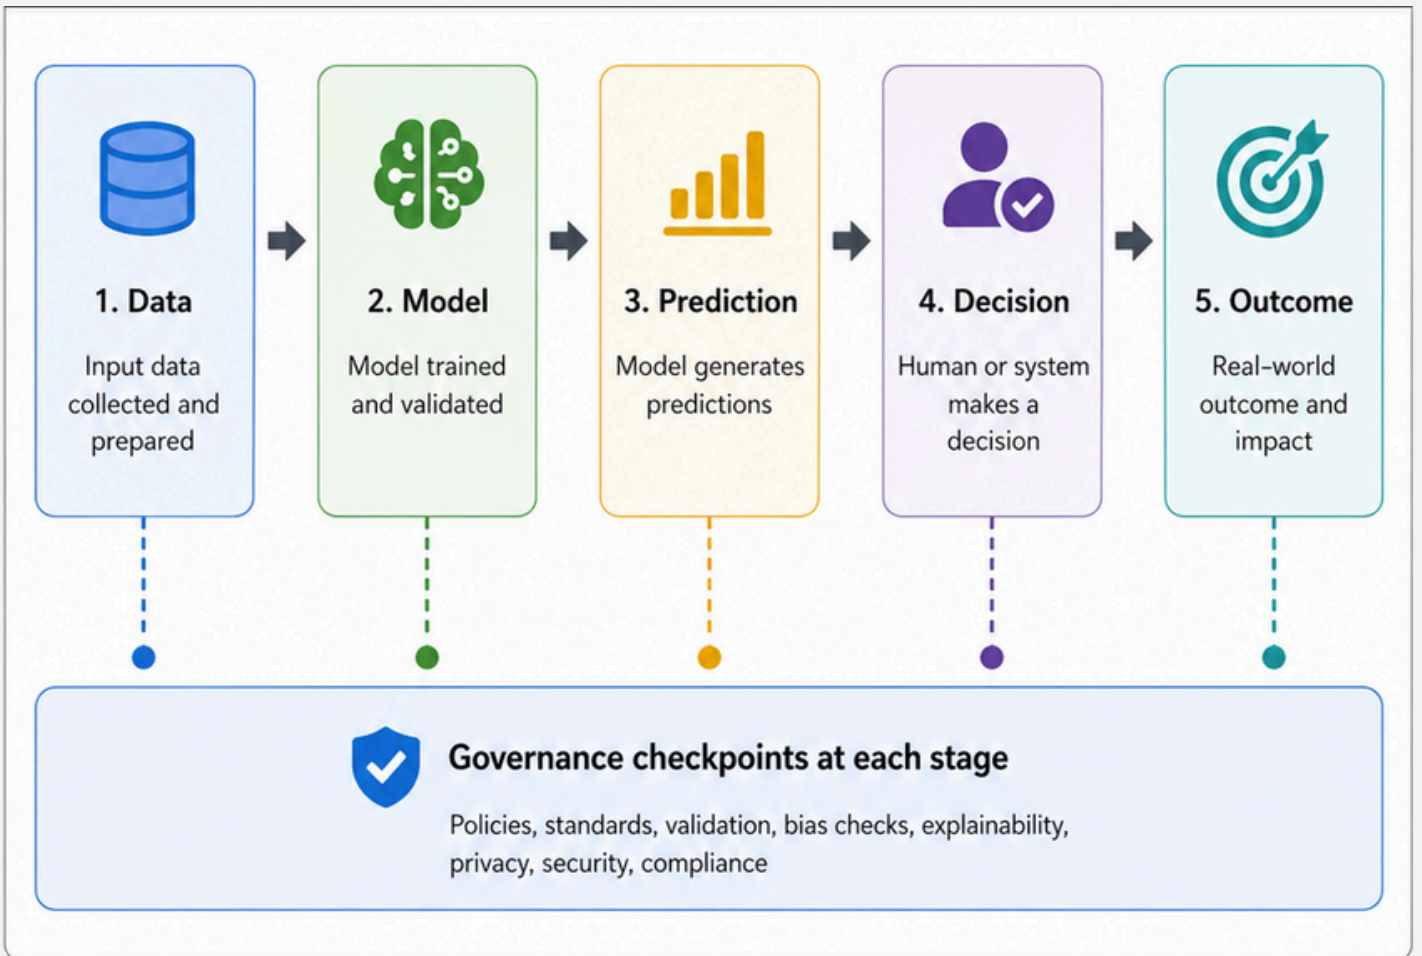
Prescriptive analytics: recommending actions to achieve desired outcomes or mitigate risks

This report focuses on the growing use of predictive AI in the corporate world. The GCC is among the fastest adopters of AI globally, with adoption shifting from a “nice-to-have” to a core decision engine. This shift is increasing board and executive accountability not only for business outcomes, but also for the quality, fairness, traceability, and governance of AI-influenced decisions. However, governance maturity remains uneven, with regulation and governance frameworks not keeping pace with adoption. As of 2024, only 39% of Fortune 100 companies reported any form of board-level AI oversight, including dedicated committees, board expertise, or ethics functions (McKinsey, 2025).

While predictive AI accelerates analysis and decision-making using larger datasets, it also introduces risks that organisations do not fully understand or control. A 2025 McKinsey survey on AI adoption in GCC countries found that 53% of respondents cited output inaccuracy as a key barrier to adoption, highlighting the gap between rapid AI deployment and organisational understanding of outputs. To ensure accountability for quality, fairness, and traceability, organisations need proactive governance. Human judgment must remain embedded, structured, and accountable at critical decision points.

At the same time, predictive AI is shifting governance from periodic review to real-time oversight, forecasting governance failures via early warning systems and automated reporting with faster auditability.

AI DECISION FLOWS



THE GOVERNANCE GAP BOARDS MUST CONFRONT

Boards operate in an increasingly uncertain world, and the risks that come with reliance on predictive AI do not sit in the server room. When AI influences business decisions, accountability sits in the boardroom.

- Opacity and limited explainability in AI-driven decisions
- Bias and unfair outcomes
- Accountability gaps across business, technology, and assurance functions
- AI tools approved as decision-support quietly becoming de facto decision-makers
- Data risks: quality, provenance, privacy, leakage, security
- Third-party and vendor risk: tools, data, outsourced models

THE REGIONAL CONTEXT: CORPORATE GOVERNANCE IN AI IN THE GCC

Governments across the region have signalled clear ambition to promote AI adoption by establishing national AI strategies, dedicated authorities, and ethics frameworks in every member state. However, as the technology evolves, regulation is still catching up. This regulatory gap creates an opportunity for organisations.

In the absence of external accountability, internal governance becomes the primary source of accountability. Good corporate governance measures implemented in AI adoption at this stage position organisations to de-risk reputational, operational, and legal exposures as regulations mature. This responsibility does not sit with technology or compliance functions alone. It sits in the boardroom.

Boards that act now and govern well may find themselves not just compliant when regulation arrives, but contributing to how it gets shaped. Regulators look to leading practice when designing frameworks. In a region moving this fast, the organisations that build robust internal AI governance today may well define what good looks like for the future.

REGIONAL REGULATORY LANDSCAPE AT A GLANCE

Country	Key Body	AI-Specific Framework	Details
UAE	UAE Council for Artificial Intelligence and Blockchain	AI Ethics Guide, National AI Strategy 2031	Most advanced regulatory environment in the region, with AI-specific legislation expected in 2026. Organisations operating in DIFC or ADGM free zones should note that additional obligations may apply.
KSA	Saudi Data and Artificial Intelligence Authority	National AI Ethics Principles, AI Adoption Framework	Strong and sector-specific. Financial services (SAMA), healthcare (SFDA), and aviation (GACA) all carry additional AI-related obligations. Boards overseeing multiple business lines should map obligations by sector.
Qatar	Ministry of Communications and Information Technology	Ethics Guidelines 2024, QCB AI Guideline 2024	Developing a framework with meaningful progress in 2024, particularly for financial institutions. The QCB AI Guideline is the most actionable board-level reference currently available.
Bahrain	Telecommunications Regulatory Authority	National AI Policy 2025	Framework established in 2025. Binding obligations remain limited internal governance carries significant weight in the interim.
Kuwait	Communication and Information Technology Regulatory Authority	National AI Strategy 2025–2028	The regulatory framework is actively being developed. Organisations should establish internal AI policies and employee guidelines now in anticipation of incoming requirements.
Oman	Ministry of Transport, Communications and Information Technology	Public Policy for Safe and Ethical Use of AI Systems 2025	Framework established in 2025 with a clear emphasis on safe and ethical deployment. Sectoral oversight bodies in finance, telecom, and health will serve as primary enforcement mechanisms as obligations develop.

In the absence of AI regulation, Personal Data and Production Laws (PDPL) apply in all GCC countries, countries promoting adoption and sandboxes for safe implementations.

WHERE SECTORS REGULATORS ARE MOVING FASTER

While national legislation continues to develop, high-risk sectoral regulators are ahead in implementing legislation. This means governance expectations vary significantly depending on what your organisation does, not just where it operates.

Financial Services

The most active sector for AI governance across the region. Central banks and financial regulators in the UAE, Saudi Arabia, Qatar, and Bahrain have issued guidance on algorithmic decision-making, fraud detection, credit scoring, and enabling technologies. For any board with exposure to lending, insurance, or investment, AI governance is already a compliance matter, not an emerging one.

Healthcare

AI used in diagnostics, medical devices, or pharmaceuticals is subject to sectoral oversight across the region, with dedicated health authorities in each member state setting expectations for safety and accountability. An AI-influenced clinical decision has immediate human consequences, which means regulatory expectations in this sector are high and tightening across the GCC.

Cybersecurity and Critical Infrastructure

Cybersecurity frameworks across the region increasingly extend to AI systems, including third-party and vendor deployments. This is an obligation that travels up the supply chain to the board. Organisations cannot delegate cybersecurity accountability for AI tools they did not build but have chosen to deploy.

What Every Board Should Be Asking

Regardless of sector or country, the following questions should be on every board agenda:

- Which regulator(s) govern our use of AI, and have we reviewed their latest guidance?
- Where is AI influencing our decisions, and who is accountable when those decisions go wrong?
- Where have we drawn the line between AI advising and AI deciding?
- Are we monitoring our models for drift, bias, and degradation over time?
- Are our vendors compliant with local guidelines? Their exposure is ours.

THE SELF GOVERNANCE IMPERATIVE

Where regulation is still emerging, compliance is the floor, not the ceiling. Where it doesn't already exist, internal governance must fill the gap. Organisations that treat regulatory absence as permission to wait are accepting governance risk without the protection of a framework to manage it. The question is not whether AI will require corporate governance. It is whether organisations will have built it before it becomes mandatory.

In practice, this does not require complex structures. Effective governance of predictive AI starts with clarity on where AI is used, who is accountable AI-influenced decisions, and where human judgement must remain embedded. Simple, well-defined guardrails are often more effective than complex frameworks that are difficult to operationalise.

CORE ELEMENTS OF YOUR CORPORATE GOVERNANCE STRATEGY FOR PREDICTIVE AI

The use of AI helps businesses become faster, more analytical, and more efficient, ultimately contributing to cost savings and value creation. However, these efficiency gains cannot replace human judgment in critical decision-making. If everyone relies on AI outputs, accountability for failures remains with the organisation, not the underlying model.

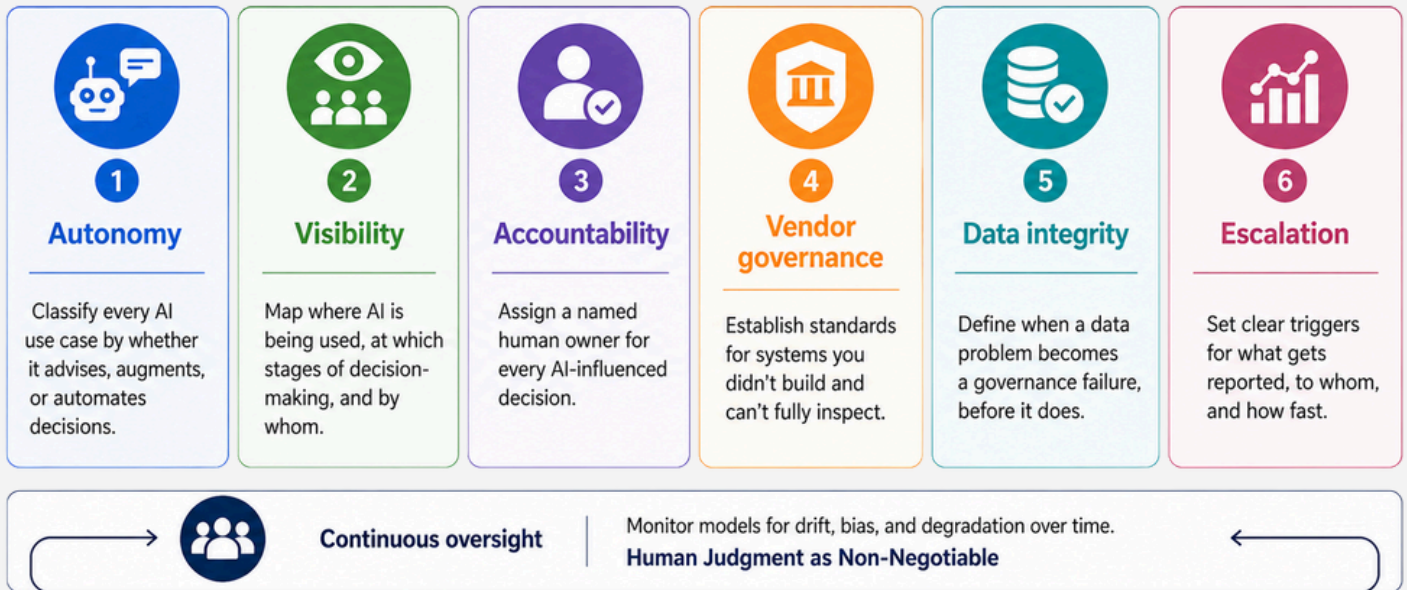
Human judgment is therefore non-negotiable and not symbolic. Boards must define mandatory human review thresholds across financial, operational, ethical, legal, and reputational decisions. Escalation triggers should be established when AI outputs conflict with policy or ethics. Employees need clear, formal guidelines for AI use, and boundaries must be set where automation is not permitted.

AI may inform decisions. It must not eliminate accountability.

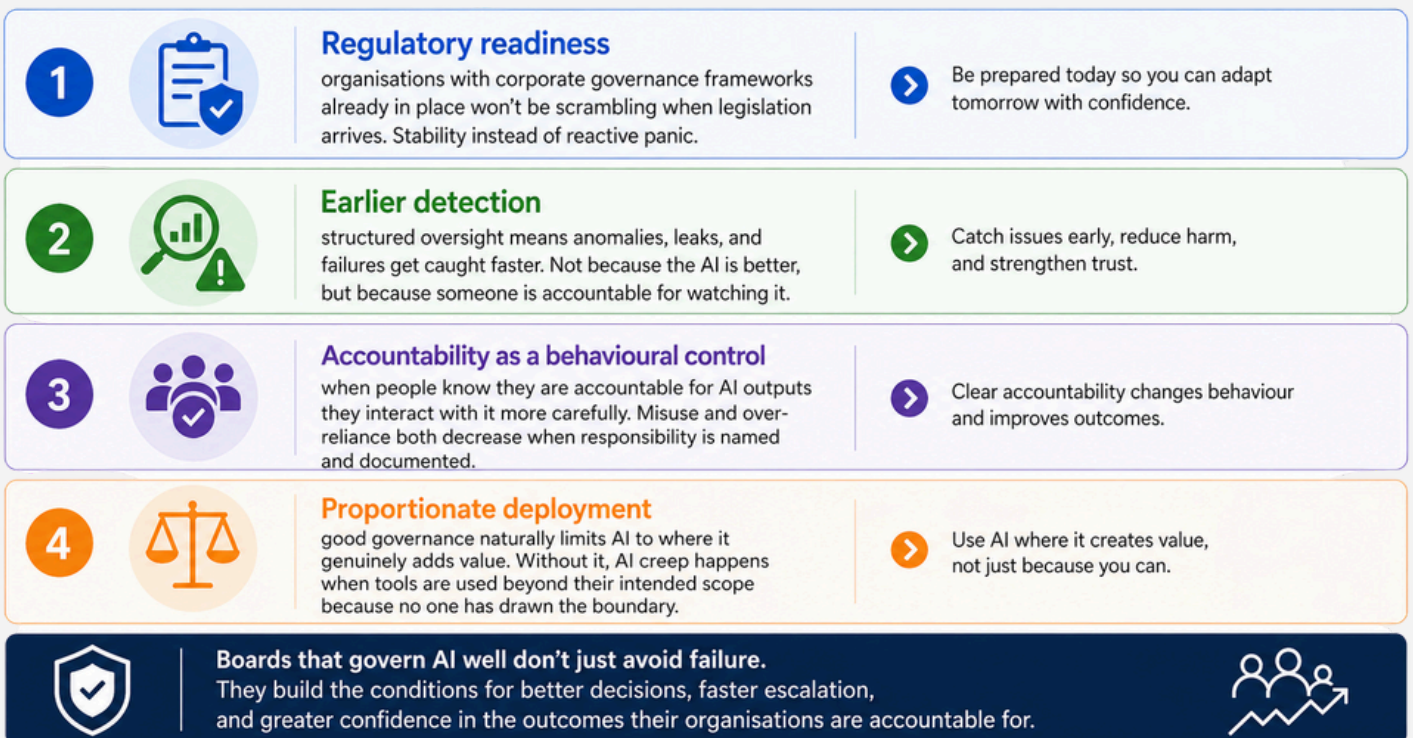
*Good Corporate Governance is Good
Business*



A ROBUST STRATEGY ADDRESSES SIX DIMENSIONS:



Governing predictive AI well is not just a risk management exercise. It is a smart business decision. Embedding good corporate governance practices in predictive AI delivers tangible benefits beyond risk mitigation:



Boards that govern AI well don't just avoid failure. They build the conditions for better decisions, faster escalation, and greater confidence in the outcomes their organisations are accountable for.

PROGRAMME PERSPECTIVES

The governance implications of predictive AI are not limited to large corporates. They extend across organisational types, each operating within different structures, priorities, and constraints.

Drawing on the Pearl Initiative’s Programmes in Governance in Tech, Governance in Family Firms, Governance in MSMEs, and Governance in Philanthropy, alongside its focus on Anti-Corruption and Diversity in Business Leadership, the following perspectives examine how predictive AI is being adopted and governed across the ecosystem.

While the applications differ, the underlying challenge remains the same: ensuring accountability, transparency, and responsible decision-making as AI becomes embedded in organisational processes. These perspectives reinforce a central point: governance is not dependent on scale or sector, but is a requirement across all organisations engaging with AI.

AI governance across organisational types

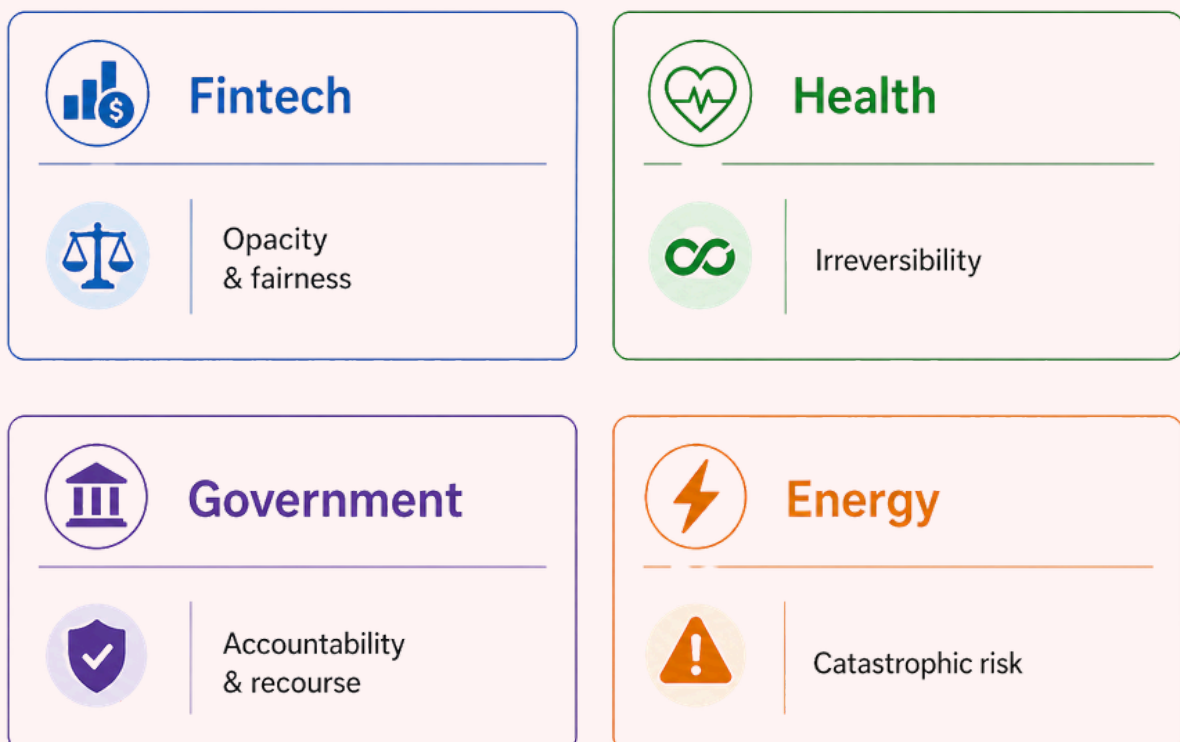
Type	Core challenge	What boards should focus on
 Tech	 Scale vs control	 Define oversight and decision boundaries
 Family	 Evolution vs Legacy	 Align AI use with family values and governance structures
 MSME	 Growth vs structure	 Introduce simple accountability mechanisms
 Philanthropy	 Capacity vs ethics	 Ensure ethical use despite constraints
 AC	 Efficiency vs accountability	 Maintain clear ownership and auditability of decisions
 DIBL	 Consistency vs inclusion	 Monitor, test, and actively mitigate bias

THE UAE'S NEXT COMPETITIVE ADVANTAGE ISN'T AI; IT'S AI GOVERNANCE

Predictive AI is already shaping decisions across the UAE's most critical sectors. The governance challenge is no longer participation, but preparedness.

The UAE has made a deliberate and ambitious bet on artificial intelligence. From the National AI Strategy 2031 to sector-specific mandates across healthcare, finance, and infrastructure, the UAE has positioned itself as a country that does not merely adopt technology but competes on it. That ambition is well-founded, but as AI systems move from automating routine tasks to making predictions on creditworthiness, patient risk, energy demand, and urban behaviour, accountability for algorithmic decision-making becomes increasingly critical.

This is not a theoretical question. It sits with CEOs and boards today, shaping decisions across the UAE's most strategic sectors. Leadership in the AI era will not be defined by speed alone, but by the institutional trust needed to sustain it.



Fintech: Speed Is the Virtue and the Vulnerability

Among the sectors examined in this report, fintech carries the highest concentration of predictive AI governance risk, and the highest stakes for getting it wrong. UAE-based fintechs and digital lenders are deploying machine learning models to make credit decisions in seconds, flag fraud in real time, and generate personalised financial product recommendations at a scale previously operationally inconceivable. The efficiency gains are tangible and the commercial logic is sound, but the governance gaps alongside them are equally real and less discussed.

The core issue is the opacity-consequence asymmetry: these systems decide who receives credit, at what rate, and under what terms, yet the models driving those decisions are often opaque to the boards responsible for oversight. In traditional credit processes, a human decision-maker and auditable rationale existed. In AI-driven systems, decisions emerge from model weights that are not interpretable in human terms, making accountability harder to trace.

The CBUAE's Open Finance Framework and ADGM's regulatory sandbox are enabling responsible innovation and deserve recognition for their approach. However, regulation is a floor, not a ceiling. Boards of fintechs and incumbent banks integrating AI into credit decisioning must go beyond compliance and ask harder questions:

- Can the organisation explain a declined decision to the affected customer?
- Can it demonstrate to regulators that its model does not systematically disadvantage any demographic group?

If the answer to either is no, the governance architecture remains incomplete, regardless of model accuracy.

Fintechs that build explainability and model governance into their operating model now will have a structural advantage as regional regulation tightens. Governance is not a constraint on innovation; it is the condition for its longevity.



Medtech & HealthTech: When the Prediction Is Personal

Predictive AI in healthcare operates at a qualitatively different level of consequence, and governance frameworks must reflect that difference. In the UAE's rapidly expanding digital health ecosystem, driven by the Dubai Health Authority's AI roadmap and the integration of AI into clinical decision support, algorithms are beginning to inform diagnoses, predict patient deterioration, flag mental health risk, and optimise surgical scheduling. Each application represents meaningful progress, but also introduces governance challenges that many boards are only beginning to address.

The central tension in medtech governance is between speed and irreversibility. A financial model with flawed outputs can be retrained, whereas a missed diagnosis or incorrect treatment recommendation cannot. This asymmetry requires a higher standard of scrutiny for AI systems, particularly where training data does not reflect the UAE's diverse patient population.

A key risk is that AI models trained primarily on Western clinical datasets may perform differently when applied to populations with varying genetic backgrounds, disease prevalence, or healthcare behaviours. The UAE's patient population is among the most diverse globally. Boards approving clinical AI deployments should require evidence of local clinical validation rather than relying on global regulatory approval as a proxy for safety and accuracy.

The opportunity is equally significant. The UAE's centralised and largely digitised health record infrastructure, young and technology-ready population, and innovation-oriented regulatory environment position it to become a global benchmark for clinically and ethically governed health AI. Institutions that prioritise governance will not only protect patients but also attract international partnerships, research collaborations, and specialist talent.

Government & Smart City: Governing the Governors

Perhaps no deployment of predictive AI in the UAE carries greater systemic significance than its integration into government services and urban infrastructure. Dubai and Abu Dhabi are among the world's most advanced smart city environments, where predictive systems already inform traffic management, utility distribution, public safety resource allocation, and the targeting of social services on a daily basis. This ambition is considerable, but it introduces a governance challenge that has no real parallel in the private sector.

At the heart of this issue is what can be described as an accountability inversion. In private sector applications of AI, market mechanisms provide a degree of recourse: dissatisfied users can switch providers, creating a competitive incentive to improve system performance and reliability. In government applications, however, such exit options are typically unavailable. Where predictive models shape decisions on social service allocation, permitting, or public safety interventions, an individual's ability to understand, challenge, or appeal those decisions is determined entirely by institutional design rather than market choice.

GOVERNANCE IN TECH

This raises the governance requirements in this space considerably above what is required in commercial contexts. Smart city deployments need not only technical oversight mechanisms but genuine public accountability infrastructure, including published algorithmic impact assessments, meaningful and accessible appeal pathways, and regular independent audits of model performance and fairness conducted by parties with no stake in the outcome. UAE government entities at the frontier of AI adoption are well-positioned to set the regional standard here, and there is real reputational and diplomatic value available to those that do. Being the first government in the region to publish a comprehensive and credible AI accountability framework is not a compliance cost; it is a statement of intent and a marker of institutional maturity.

Accountability in government AI must be deliberately designed into the system architecture. It will not emerge as a byproduct of good intentions or technical sophistication alone.

”

Energy & CleanTech: The Long Game

The energy sector may operate on a longer governance horizon than other verticals, but this should not be read as an absence of near-term urgency. Across ADNOC's AI-driven reservoir management, DEWA's grid optimisation, and the renewables ecosystem around Masdar, artificial intelligence is becoming foundational to managing critical national assets and delivering net-zero commitments at COP28.

The governance challenge in energy AI is defined less by the likelihood of failure and more by its consequences. Predictive maintenance systems are valuable but depend on not missing critical signals. When failures occur in high-pressure infrastructure or large-scale grids, the operational and environmental impact is far greater than in other sectors. Governance frameworks must reflect this asymmetry, with rigorous stress-testing, clear override procedures, and human-in-the-loop requirements for high-risk decisions.

A longer-term governance issue is also emerging. As ESG scrutiny intensifies and the UAE's COP28 legacy demands credible climate action, predictive AI will underpin emissions measurement, reporting, and verification. The integrity, auditability, and transparency of this data are governance issues as much as technical ones, and organisations that recognise this will be better positioned as scrutiny increases.

The Role of the Board

Across these four verticals, a common thread emerges: the gap is not primarily technical, but one of leadership. Boards and executive teams have often treated AI as the domain of the CTO or data science function, but this approach is no longer sustainable.

Predictive AI is making decisions with legal, ethical, reputational, and sometimes physical consequences, making it a board-level concern by definition. The question is not whether boards have the technical expertise to audit models, but whether they are asking the right questions and ensuring accountability is appropriately allocated.

Boards should focus on:

- **Where predictive AI is used across the organisation**
- **Who is accountable for AI-influenced decisions**
- **Whether governance structures ensure clear escalation and oversight**
- **Where human judgement must remain embedded**

The UAE's ambition in artificial intelligence is clear, supported by infrastructure, regulation, and political will. What remains in development is the governance architecture that will determine whether this ambition delivers sustainable outcomes. Organisations and boards that invest in this now will not be slowed by it, but differentiated by it.

These examples point to a broader shift: as predictive AI becomes embedded in critical systems, governance must evolve to ensure that speed and scale do not come at the expense of accountability.

THE EVOLVING ROLE OF AI IN FAMILY BUSINESS GOVERNANCE

Family-owned enterprises remain a cornerstone of global and regional economies, contributing significantly to GDP, employment, and long-term economic stability. The 2025 EY and University of St. Gallen Global 500 Family Business Index reports that the 500 largest family businesses generate aggregate revenues of US\$ 8.8 trillion.

In the GCC, family businesses represent more than 50% of the private sector economy and play a critical role in diversification, national development agendas, and long-term economic visions.

With an estimated \$1 trillion in assets set to transfer to the next generation over the next decade, family enterprises sit at a critical intersection of legacy and transformation, making governance more critical than ever. They face unique challenges linked to ownership structures and family dynamics, making governance essential to balance complexity and ensure continuity. Strong governance supports:

- Alignment and accountability across family members, management, and employees
- Clear communication and reduced conflict
- More consistent long-term value creation

In this context, Artificial Intelligence (AI) is emerging not only as an efficiency tool but as a governance enabler. Its value in family enterprises lies in improving decision consistency, particularly in:



AI also strengthens governance by:

- Grounding decisions in data rather than perception
- Improving communication and objectivity in sensitive discussions
- Supporting forward-looking strategic planning
- Enhancing alignment while preserving human judgement and family values

Finally, AI introduces greater objectivity by analysing financial, market, and operational data, reducing bias and supporting fairer, evidence-based decisions that limit disputes.

GOVERNANCE IN FAMILY FIRMS

Succession planning remains one of the most critical governance challenges for family businesses, often determining long-term success or decline. AI presents an opportunity to enhance this process through a data-driven approach, reducing subjectivity in decision-making. Through predictive analytics, AI can assess historical performance, leadership capability, and future business needs to support succession planning.

At an operational level, AI can automate routine processes, optimise resource allocation, and generate forward-looking insights, improving resilience and responsiveness to market change. This drives efficiency while allowing leadership to focus on higher-value priorities such as growth, diversification, and innovation. However, these benefits require robust governance frameworks to manage risks and ensure alignment with business objectives and values.

Key governance considerations include:

- Data privacy and protection
- Transparency in decision-making
- Fairness and absence of bias in outcomes

These are essential to maintaining trust and accountability within family business structures. As AI becomes more embedded, organisations must implement safeguards and clear ethical frameworks to mitigate risks such as algorithmic bias, data misuse, and oversight gaps. Without this, decision quality and stakeholder confidence may be undermined.

Boards and leadership teams must therefore play an active role in:

- Setting clear policies for AI use
- Establishing monitoring and oversight mechanisms
- Embedding accountability across the organisation
- Ensuring alignment with family values and legacy

At the same time, AI adoption must be balanced with the preservation of culture, legacy, and values central to family enterprises. This requires structured change management, communication, and training to align stakeholders.










Ultimately, AI presents both an opportunity and responsibility. When guided by strong governance, it can improve decision-making, strengthen succession planning, and enhance resilience while preserving the identity of family businesses. Its successful integration will determine how effectively these enterprises adapt, compete, and sustain value across generations.

GOVERNANCE IN FAMILY FIRMS

The UAE's ambition in AI is clear, supported by infrastructure, regulation, and political will. What remains in development is the governance architecture that will determine whether this ambition delivers sustainable long-term outcomes. Organisations and boards that invest in this architecture now will not be constrained by it; instead, it will become a differentiating factor.

These examples reflect a broader shift. As predictive AI becomes embedded in critical systems, governance must evolve in parallel to ensure that speed and scale do not come at the expense of accountability.

Governing AI in Family Businesses: Balancing Structure and Legacy

AI-driven side	Family-driven side	Governance focus
 Algorithmic decisions	 Relationship-based judgement	 Define when decisions follow the model and when human override is required
 Objective outputs	 Values-led judgement	 Ensure decisions are explainable and aligned with family values
 Standardised decisions	 Legacy-driven exceptions	 Establish clear rules for exceptions, with accountability and documentation

PREDICTIVE AI IN MSMEs: MOVING FAST WITHOUT LOSING CONTROL

Micro, Small, and Medium Enterprises (MSMEs) are the backbone of any economy, driving employment, innovation, and growth. In the GCC, MSMEs account for over 90% of businesses, mostly family-run, high-growth, and digitising fast. Governments across the region are actively promoting entrepreneurship, making it easier for both locals and expatriates to establish and scale businesses. As digital adoption accelerates across the GCC, many MSMEs are integrating AI tools earlier in their growth journeys than ever before.

These businesses are emerging rapidly, some seeking access to regional capital, some seeking a stable market, some seeking a share of a robust and growing economy. All are focused on survival and growth. In this race, governance often ends up on the back burner.

In today's business landscape, predictive AI allows MSMEs to be more competitive, faster, by predicting success and failure scenarios, anticipating market response, and managing volatility. But without a governance structure around how AI is used, the same tools that accelerate growth can accelerate failure. However, Gartner predicts that 60% of AI projects will be abandoned through 2026 if they lack the governance and data quality to sustain them. For example, McDonald's ended its 3-year AI-driven drive-thru ordering pilot after repeated errors highlighted the risks of deploying AI without sufficient oversight. What was intended to streamline operations instead introduced operational friction and reputational risk. While at a larger scale, this highlights a broader risk for growing businesses: without governance, AI can just as quickly undermine operations as it can enhance them.

Predictive AI needs to be used smartly and needs to be derisked

1. Chasing trends that don't serve core business needs

It's easy for MSMEs to pivot toward what's new or popular, particularly in fast-moving spaces like AI. However, tools should respond to clearly defined business needs, not external hype. Adoption without purpose can dilute focus and resources.

2. Adopting tools solely to signal investor readiness

Implementing systems or technologies purely to appear "advanced" can backfire. Investors value clarity, efficiency, and intentionality. Overbuilt or unused systems often signal weak decision-making rather than preparedness.

3. Cutting costs at the expense of quality and risk management

Choosing tools based only on price can introduce operational, security, and reliability risks. Short-term savings may result in higher long-term costs, particularly if systems fail or require replacement.

4. Using tools without fully understanding their implications

Signing up for platforms without reviewing terms, data policies, or access conditions is a governance risk. Data ownership, storage, and usage rights are not administrative details; they are strategic assets. Your data is part of your competitive advantage and must be treated as such.

GOVERNANCE IN MSMEs

Start with low-risk, high-value applications before expanding into critical decisions. For instance, an MSME using predictive tools to forecast demand without reliable historical data may overproduce inventory or misallocate resources. Without oversight, such decisions can quickly strain cash flow and operations.

In practice, this does not require complex systems. For MSMEs, effective AI governance often starts with a few clear principles: clarity on purpose, accountability for decisions, and ongoing oversight of outcomes. These simple foundations can significantly reduce risk while preserving the agility that defines smaller businesses.

Why governance matters for MSMEs adopting AI

Governance should be simple and right-sized. Strong governance in MSMEs is not a constraint on agility; it is a signal of maturity and a foundation for sustainable growth. Particularly in the context of AI adoption, governance enables businesses to move fast without losing control.

- Signals founder maturity and discipline to investors and external stakeholders
- Reinforces human accountability in AI-assisted decision-making particularly for high-stakes or irreversible decisions
- Reduces regulatory risk and avoids costly pivots as frameworks evolve
- Provides operational stability and consistency in decision-making
- Strengthens trust with partners, clients, and financial institutions
- Improves data integrity, leading to more reliable AI outputs as poor data directly leads to flawed predictions and decisions, for example: inaccurate sales or customer data leading to flawed demand forecasts
- Reduces over-reliance on founders by introducing structure and oversight
- Ensures ongoing monitoring of AI outputs to detect errors, bias, or drift over time

The role of boards in MSMEs

While many MSMEs operate with lean, founder-led structures, introducing even a light governance layer, such as an advisory or formal board, can significantly strengthen decision-making even in an informal or advisory capacity.

Operating solo allows for speed, but it also concentrates risk. Boards provide a structured space for challenge, oversight, and reflection.



- Act as a sounding board for complex or high-stakes decisions
- Introduce external perspectives and reduce blind spots
- Support more disciplined, long-term thinking
- Help balance speed with accountability

GOVERNANCE IN MSMEs

For MSMEs, the challenge is not whether to adopt AI, but how to do so responsibly. Unlike larger organisations, they often lack the buffers to absorb failure, making early decisions around governance even more critical. Putting in place simple, practical guardrails from the outset allows businesses to experiment and innovate while maintaining control.

As AI continues to shape competitive advantage across the region, MSMEs that invest early in governance will not only reduce risk, but position themselves as more resilient, credible, and investment-ready in the long term.

Governing AI in MSMEs: Moving Fast Without Losing Control

Growth-driven side	Risk side	Governance focus
 Moving fast to compete	 Adopting tools without purpose	 Define where AI adds real business value
 Lean, founder-led decisions	 Concentrated risk and blind spots	 Introduce basic oversight and challenge mechanisms
 Cost-driven choices	 Data, security, and reliability risks	 Set non-negotiable standards for tools and data use

PURPOSE, TRUST AND ACCOUNTABILITY: GOVERNING PREDICTIVE AI IN PHILANTHROPY

The philanthropic ecosystem is robust, spanning individual donors, institutional funders, grantees, impact investors, and social enterprises, addressing issues from humanitarian crisis and medical innovation to poverty, economic empowerment, and refugee support. The ecosystem catering to diverse needs is constantly facing high demand and limited resources. AI offers speed and agility to the sector, but adoption has remained mostly surface-level, generative rather than predictive. As the pace of adoption picks up, organisational governance of AI systems becomes essential.

Philanthropic organisations make decisions about who receives support, which causes get funded, which communities get prioritised. When AI influences those decisions, the accountability stakes are uniquely reputational and ethical in a way that's distinct from commercial organisations. A right decision made faster with the use of ai can mean lives saved and a wrong decision with ai bias can mean lives lost.

In the GCC, the philanthropic ecosystem is robust, with cultural and faith-based giving taking place at a national, corporate, foundation, family, and individual level. Beneficiaries span across the world. As one of the fastest adopters of AI globally, the GCC is seeing increased use of AI being introduced into the philanthropic ecosystem, influencing decision-making with little to no structural organisational governance in place.

The Investment Gap

Within the ecosystem, funding is directed primarily towards mission delivery. Securing operational funding is already a challenge, and investment in technology infrastructure and governance capacity is even more limited. Research by the Center for Effective Philanthropy in 2025 found that nearly 90% of foundations provide no AI implementation support to their grantees. This gap extends beyond funding to governance capacity itself.

This is particularly critical as many organisations are adopting AI tools without dedicated technical or governance capacity, increasing the risk of misapplication and unintended consequences. Without this support, the sector cannot govern AI responsibly. Funders who expect responsible AI adoption from their grantees have an obligation to support the governance capacity that makes it possible.

Benefits of Predictive AI in the Philanthropic Ecosystem

- **Crisis preparedness:** Early prediction of natural disasters, economic crises, and geopolitical shifts allows the ecosystem to mobilise resources before a crisis hits rather than in response to it.
- **Beneficiary identification:** Data-driven prioritisation of who needs support first, reducing reliance on incomplete or anecdotal information.
- **Donor management:** Relationship management informed by giving patterns and engagement data, avoiding both under and over outreach.
- **Deployment of capital:** Analysis of where and how much funding can be deployed for the greatest impact, supporting more defensible allocation decisions.
- **Proactive risk mitigation:** Early identification of operational, reputational, and programmatic risks for both funders and implementing organisations.
- **Decision support and triage:** AI-assisted prioritisation of competing needs, with human judgement retained for final decisions.
- **Optimised resource allocation:** Ensuring limited resources reach the right places at the right time, reducing waste and improving outcomes.

Predictive AI in Action

1. Crisis Response at Scale: World Food Programme

In 2023, during the earthquake that struck Türkiye and Syria, the WFP deployed AI for machine learning-based damage assessment, near real-time monitoring, and enhanced situational awareness. A team of five individuals was able to assess approximately 600,000 buildings with an accuracy of over 81%, impacting over five million people. The data gathered empowered decision makers with the necessary information to plan and execute relief efforts, ensuring assistance reached those who needed it most.

The governance observation: speed and scale at this level are only trustworthy when there is clear human accountability for how outputs are acted upon. The 19% inaccuracy rate in a crisis context is not a technical footnote. It is a governance question about who reviews, who decides, and who is responsible for what the model missed.

2. Anticipating Disaster Before It Strikes: Mercy Corps

When war broke out in Sudan in 2023, Mercy Corps used ten years of satellite imagery showing crop health patterns to compare real-time conditions with historical data. AI gave their team the ability to identify the most vulnerable areas and deploy resources before disaster struck rather than in response to it.

The governance observation: predictive tools of this kind require defined accountability for how vulnerability data is collected, stored, and acted upon. Clear human oversight is essential at the point where a prediction becomes a resource decision.

Risks: Why Organisational Governance is Non-Negotiable

- **Bias:** AI outputs reflect the data they are trained on. In a sector serving marginalised communities, biased models risk compounding existing inequities.
- **Inconsistent or inaccurate data:** Poor data quality produces unreliable outputs, with consequences that in this sector can directly affect vulnerable populations.
- **Accountability gaps:** When AI influences funding or beneficiary decisions, it must be clear who is responsible for those outcomes.
- **Talent development:** The sector cannot govern what it does not understand. Building AI literacy across staff and boards is a governance prerequisite.
- **Data sensitivity:** Both beneficiary and funder data carry significant privacy obligations that AI systems must be designed to protect.
- **Funder bias in data and outputs:** AI systems trained on historical funding patterns risk perpetuating the same biases those patterns contain.
- **Vendor reliance:** Most organisations will buy rather than build. Governance frameworks must extend to the tools and vendors they depend on.
- **Regional data gaps:** In the GCC, limited availability of localised data means AI systems may rely on datasets that do not reflect regional realities, undermining both accuracy and relevance of outputs.
- **Overstatement of impact:** AI-generated impact projections must be validated against real outcomes. In a sector built on donor trust, inaccurate reporting is a governance failure not just a measurement error.

With the impact of risks being so high, human judgment and accountability remain critical. Governance measures to ensure accountability and ethical use of AI within organisations are not optional they are a prerequisite for responsible adoption.













Philanthropic boards govern for mission, funders, beneficiaries, and the do no harm principle simultaneously. That is a governance structure that commercial boards do not face. AI adds complexity to each of those obligations, making board-level oversight not just relevant but essential.

Unlike commercial contexts, beneficiaries often have limited visibility into or ability to challenge decisions, making governance not just a structural requirement, but an ethical obligation.

What makes the philanthropic ecosystem uniquely complex is that governance failures do not stay contained. An individual donor's giving decisions informed by AI, a foundation's grant screening informed by AI, a grantee's beneficiary identification informed by AI. Each decision point connects. Governance at one level affects outcomes at every other level.

GOVERNANCE IN PHILANTHROPY

Governing Across the Philanthropic Value Chain

Stage	Where AI is used	Governance focus
 Donor / Funder	 Prioritisation of causes and funding strategy	 Ensure transparency in decision criteria and alignment with mission
 Foundation / Organisation	 Allocation of funds and programme design	 Maintain clear accountability and oversight of AI-informed decisions
 Grantee / Implementer	 Targeting of beneficiaries and delivery decisions	 Validate data quality and monitor for bias in targeting
 Beneficiary / Community	 Impact measurement and outcome tracking	 Ensure ethical use, accuracy of outcomes, and human review

The sector cannot afford to be a late adopter. With a unique positioning at the intersection of academia, government, and business, the philanthropic ecosystem has the opportunity to set the tone on organisational governance of predictive AI. With many Fortune 100 companies still struggling with implementation, the philanthropic ecosystem, as the most vulnerable yet most diverse actor in this space, is uniquely placed to lead. Boards need to be proactive, funders need to support adoption, and the third sector needs governance frameworks now before it falls behind and plays catch-up later.

In this context, governance is not only about effectiveness, but responsibility. Ensuring that AI-enabled decisions do not unintentionally exclude, misdirect, or harm the communities they are intended to serve.

ANTI-CORRUPTION AND COMPLIANCE IN THE AGE OF ARTIFICIAL INTELLIGENCE

Corporate governance has long been the foundation of responsible business conduct. At its core, governance establishes the structures, rules, and oversight mechanisms that ensure organizations act ethically, comply with regulations, and remain accountable to all their stakeholders. Effective governance embeds strong anti-corruption frameworks, robust compliance systems, and a culture of ethics into the organisation.

The emergence of Artificial Intelligence (AI) is transforming how organisations operate and make decisions. Businesses are increasingly leveraging AI to strengthen their corporate governance frameworks through data driven, real time insights and predictive decision making. This in turn enhances risk management, reduces regulatory risks, and drives greater operational efficiency across organisations.

As AI becomes increasingly integrated into business processes effective governance requires a holistic approach that integrates ethics, anti-corruption systems, compliance mechanisms, and technological accountability into the design, deployment, and oversight of AI-enabled systems.

Integrating AI into governance systems is not just about technology. It needs to be embedded within broader organisational systems of ethics, anti-corruption, and compliance. This integration ensures that technological innovation aligns with the organization's values, legal obligations, and societal responsibilities.

While AI offers immense opportunities to enhance efficiency, accountability and transparency it also introduces new ethical, governance, and compliance challenges.

Benefits of AI:

- **Enhanced decision-making:** Strengthen governance by analysing historical data, market trends, and financial indicators to generate actionable insights. This enables more informed, timely, and evidence-based decisions.
- **Improving board effectiveness:** AI-driven dashboards and analytics provide boards with clear, real-time visibility into performance, risks, and regulatory developments.
- **Strengthening risk management:** Risk management is a critical component of corporate governance. AI can significantly enhance risk management by identifying risks and patterns that indicate potential issues.
- **Automating compliance processes:** Automate routine compliance tasks reducing operational burden, improving accuracy, and minimising risk of human error.
- **Enhancing reporting:** Streamline reporting by collecting, analysing, and structuring large volumes of data.

Challenges of AI:

- **Ethical and Data Governance Concerns**

The use of AI in governance and compliance raises concerns around data privacy, transparency and bias necessitating the need to ensure that AI systems are explainable, fair, and supported by robust data governance frameworks to maintain stakeholder trust.

- **Accountability and Responsibility Gaps**

AI-driven decision-making can blur lines of accountability. Determining responsibility across developers, management, and boards remains complex, necessitating clear governance structures and defined oversight roles.

- **Integration and Adoption Challenges**

Embedding AI into existing governance and compliance frameworks requires significant investment in infrastructure, workforce upskilling, and well-defined implementation strategies.

- **Regulatory Uncertainty and Compliance Risks**

The evolving global regulatory landscape creates uncertainty for organisations. Companies need to remain agile, continuously monitor developments, and ensure compliance with emerging laws and standards across jurisdictions.

Responsible AI Adoption: Governance, Ethics, and Oversight

As organisations increasingly integrate artificial intelligence into their operations, it is imperative that adoption is guided by strong systems, ethical principles, and robust oversight mechanisms. While AI offers significant opportunities to enhance efficiency and decision-making, it also introduces risks. A structured and well-governed approach is therefore critical to ensure that AI supports responsible, transparent, and value-driven outcomes.

- **Define clear purpose and capability requirements:** Organisations should begin with a clear understanding of the problems AI is intended to solve and the internal capabilities required to manage it effectively.
- **Embed ethics and accountability:** Ethical considerations must be integrated from the design stage of AI systems. Organisations should foster a culture of critical thinking that looks at validating AI outputs to ensure right interpretation of outputs, identify limitations, and use AI responsibly.
- **Strengthen monitoring, data governance, and risk controls:** The effectiveness of AI depends on the quality and integrity of underlying data. Organisations must implement robust data governance frameworks, ensure compliance and embed monitoring and controls to mitigate, misuse and unethical outcomes.

ANTI-CORRUPTION

- **Ensure leadership oversight and organisational alignment:** Strong leadership commitment and board-level oversight are critical for responsible AI adoption. This includes building relevant expertise and including board members with relevant expertise.













Conclusion

The integration of artificial intelligence into corporate governance frameworks offers immense opportunities for improved efficiencies and innovation through enhanced decision making, risk management and compliance processes. However, realising these benefits requires a deliberate and responsible approach that includes ethics, transparency, accountability, and data protection.

Ensuring that AI initiatives are aligned with long-term corporate objectives and is supported through clear policies, regulatory alignment, and strong oversight mechanisms is essential to driving sustainable value creation, trust and competitive advantage.

AI in corporate governance is not merely a technological advancement but a strategic imperative to strengthen institutional trust and achieve long-term resilience and success in an increasingly complex business environment.

AI in Compliance: From Automation to Accountability

AI Enables	Risk	Governance Must Ensure
 Faster decisions	 Blurred accountability	 Clear ownership of decisions
 Automated compliance	 Reduced transparency	 Explainability of outputs
 Real-time monitoring	 Over-reliance on systems	 Human oversight at key points
 Data-driven insights	 Bias and data issues	 Strong data governance

USING AI TO STRENGTHEN DIVERSITY AND INCLUSION IN ORGANISATIONS

In today's rapidly evolving global landscape, transforming operations and workplace culture to prioritise diversity and inclusion is not just a moral imperative; it is a strategic necessity to drive growth, innovation, agility, and long-term resilience. As businesses confront numerous challenges, including growing economic pressures, changing social norms, and rapid advancements in technology and artificial intelligence, inclusivity has shifted from being a desirable goal to an essential strategy for thriving.

Building a workplace that values unique perspectives and embraces differences can lead to enhanced collaboration, creativity, and decision-making within the business. Diverse teams bring in a wider range of ideas and solutions, helping businesses navigate uncertainties and adapt to rapid changes with greater agility.

However, managing D&I in a consistent and measurable way remains a challenge. While many organisations have policies and targets in place, translating these into meaningful outcomes often requires more structured and data-driven approaches. In this context, artificial intelligence (AI) has begun to play a significant role in supporting how organisations manage D&I functions and is being used to enhance decision-making, strategic planning, and bring greater consistency to processes that were traditionally subjective. AI can analyse large volumes of data, identify patterns, and provide structured insights, enabling organisations to move from a reactive to a more proactive approach in managing diversity and inclusion.

Key Applications of AI in DEI Accountability:

A significant benefit of AI in managing diversity and inclusion is its ability to identify and address disparities at scale through data-driven insights. AI can analyse patterns in hiring, progression, engagement, and access to opportunities across different groups, making gaps more visible and enabling more targeted and equitable interventions. AI can support organisations in developing a deeper understanding of workplace culture, employee engagement, and overall experience. By analysing employee inputs such as feedback, surveys, and interaction patterns, AI can help identify trends related to inclusion, sense of belonging, and well-being across employees, highlighting disparities that may otherwise not be immediately visible.

AI can also be used in workforce analytics, where it can be used to analyse employee data and transform it into actionable insights for optimising productivity, engagement, and operational resilience. This allows for proactive and targeted intervention strategies that address specific challenges.

DIVERSITY IN BUSINESS LEADERSHIP

AI is also being widely adopted in recruitment and talent management processes to screen resumes, analyse and match candidates to roles, and provide data-driven insights, improving hiring speed. From a D&I perspective, this can help minimise conscious and unconscious bias and create consistent evaluation criteria in application screening and hiring decisions.

AI is increasingly being used in employee learning and development to analyse employee performance, skills, and learning patterns to recommend personalised development and engagement pathways. This can support ensuring equitable access to growth, progression, and opportunities.

Leadership plays a central role in shaping inclusive workplaces, particularly in ensuring that diversity translates into equitable participation, opportunity, and outcomes. AI can support this process by providing clearer visibility into workplace culture, team dynamics, gaps, and biases, such that diverse and inclusive perspectives are actively considered in decision-making.

Key benefits of using AI in Workplace D&I management:

- **Greater consistency:** AI reduces variability in decision-making by applying standardised criteria.
- **Improved visibility:** Organisations gain deeper insights into workforce composition and employee experience.
- **Scalability:** AI enables D&I initiatives to be applied consistently across large and complex organisations.
- **Proactive management:** Predictive capabilities allow organisations to address issues before they escalate.

Key challenges and limitations

While AI offers clear advantages in supporting diversity and inclusion efforts, its effectiveness is closely tied to how it is designed, implemented, and governed. Without the right safeguards, there is a risk that AI may reinforce existing gaps rather than address them. In practice a balanced approach where AI functions as a support tool along with human oversight is essential to ensure fairness and accountability in diversity and inclusion management in the workplace. Key challenges to consider include:

- **Data bias:** AI systems are trained on historical data, and if that data reflects existing biases, those patterns can be replicated in outcomes. Ensuring that datasets are representative and that algorithms are regularly tested and refined is critical to mitigating this risk.
- **Limited ability to capture complexity:** Diversity and inclusion involve factors such as culture, context, and lived experiences which are not always quantifiable and hence may not be visible in datasets leading to exclusion or misinterpretation without the necessary contextual understanding.

DIVERSITY IN BUSINESS LEADERSHIP

- **Employee trust and transparency:** The use of AI in managing employee data raises important concerns around transparency in how the data is being collected, analysed, and used in decision-making. Concerns around privacy and potential misuse of data can undermine broader inclusion efforts.
- **Over-reliance on technology:** There is a risk of over-reliance on AI-generated metrics and insights; while AI can inform decision-making, it cannot substitute the human elements of judgment, context, and empathy.

In practice, the value of AI lies in its ability to generate deeper and timely information and insights which could otherwise get missed. However, these insights only create value when they are translated into clear actions, decisions, and organisational practices. This in turn requires a thoughtful integration of AI within existing governance structures. This includes clear accountability, ongoing monitoring, transparent communication, and a continued emphasis on human oversight. When combined with strong leadership and organisational commitment, AI can support more effective and equitable D&I outcomes, but it cannot deliver these outcomes on its own.

Governing AI in Diversity & Inclusion

 AI Enables	 Risk	 Governance Must Ensure
 Standardised decision-making	 Reinforcing existing bias	 Regular testing and mitigation of bias
 Data-driven insights	 Loss of context and lived experience	 Human interpretation alongside AI outputs
 Scalable D&I processes	 Over-reliance on metrics	 Balanced use of qualitative and quantitative inputs
 Workforce analytics and monitoring	 Employee trust and privacy concerns	 Transparent data use and clear communication

INDUSTRY PERSPECTIVES

CORPORATE GOVERNANCE IN THE AGE OF PREDICTIVE AI


The following section brings together perspectives from the Pearl Initiative's corporate partners across the region, each sharing how predictive AI is being adopted and governed within their organisations. The editorial section of this report establishes why predictive AI is a board-level governance responsibility. The perspectives that follow focus on how this is being implemented in practice. Drawn from organisations across financial services, luxury retail, logistics, construction, and professional advisory, these contributions do not present a single governance model. Instead, they offer practical accounts of governance as it is being built in real organisations under real constraints.

Read together, they reflect a journey familiar to many boards. It begins with defining the governance architecture required, moves through the practical realities of implementation across different organisational structures and risk profiles, and concludes with a reflection on what this shift means for governance as a discipline, not just a control function.

A consistent question runs through each contribution: as predictive AI moves from supporting decisions to shaping them, where does accountability sit, and how is it made operational? Each organisation responds from its own context, but together they point to shared principles. These principles are increasingly important to define now, before regulation defines them externally.



KPMG



The starting point for any organisation governing predictive AI is architecture: what structures, controls, and accountability mechanisms need to exist, and how should they be designed? KPMG's contribution establishes this foundation, offering a practical framework for classifying AI decisions by autonomy level, embedding human oversight across the lifecycle, and building governance that is measurable rather than merely stated. It is a rigorous opening position, and a useful reference point for everything that follows.

Corporate Governance in the Age of Predictive AI: From Insight to Foresight **by Jeyapriya Partiban**

Boards are used to governing decisions. What is new is, governing decision-makers that learn, adapt, and operate at machine speed.

Predictive AI is now embedded in core business judgments: credit decisions, fraud detection, risk triage, demand forecasting, preventive maintenance, churn prediction, pricing, and workforce planning. The shift is subtle but profound. We are moving from insight (analytics that explains what happened) to foresight (systems that recommend what to do next, and sometimes do it automatically). When that happens, governance can no longer sit only around policies and committees. It must extend into model design, data provenance, deployment controls, and ongoing monitoring.

This is exactly where many organizations feel the pressure. Technology is developing and becoming readily available, while laws and regulation are still evolving and public opinion is being shaped in real time. The governance challenge is to move quickly without creating blind spots that can potentially lead to regulatory, reputational, or operational failures. Global standards are emerging to shape AI governance - the EU AI Act introduces risk-based classification and transparency obligations; ISO/ IEC 42001 sets out AI Management System requirements, and NIST's AI Risk Management Framework emphasizes trustworthiness and resilience. Boards should anticipate and consider these frameworks as potential benchmarks for compliance and ethical practice.

The governance challenge is to move quickly without creating blind spots that can potentially lead to regulatory, reputational, or operational failures particularly where AI systems evolve post-deployment in ways that were not explicitly approved or fully understood at design time.

The board question that matters: Is AI advising, or deciding?


A practical governance move is to classify every AI use case based on how much autonomy it has:

- **Automated:** the model makes the decision (for example, real-time transaction blocking).
- **Augmented:** the model recommends, humans approve (for example, underwriting exception approvals).
- **Human-supported:** the model flags and prioritizes, humans decide (for example, risk-based audit planning).

This classification is not academic. It determines the strength of controls required, the level of explainability needed, and where accountability must sit. KPMG’s Trusted AI guidance explicitly calls out the need to determine whether model decisions should be automated, augmented, or human supported, and to embed Human-in-the-Loop controls accordingly.

Governance needs an operating model, not a policy binder.



 **Continuous governance, continuous improvement.**
 Each stage informs the next to ensure responsible, transparent and effective AI systems.

Evidence for common failure modes are already emerging. AI solutions approved as “decision-support” tools are gradually being used as de-facto decision-makers; model performance is degrading while accountability sits ambiguously between IT and the business teams; or transparency commitments are made externally that cannot be substantiated internally. These are predominantly failures in governance, not technology implementation.

Many organizations start with an AI policy. Mature organizations build an AI governance operating model: decision rights, lifecycle checkpoints, risk controls, and performance oversight.

A proven construct is a **centralized AI Governance Board/ Committee** comprising AI specialists, legal and compliance leaders, business owners, and a representative from enterprise architecture oversight. The remit should include setting AI policies and standards, ensuring regulatory compliance, overseeing ethical considerations, approving AI projects, and conducting risk assessments across the portfolio.

At the same time, governance cannot become a bottleneck. This is why a **hub-and-spoke model** works well in practice: the central governance body defines the guardrails, while functions retain ownership of use cases and data, implementing within those guardrails. This model is designed to foster collaboration, ensure consistency and compliance, and promote data ownership at the departmental level. In governance terms, this is how you scale control without centralizing every decision.

The Trust Agenda: Transparency, Accountability, and Responsible Data Use

Predictive AI governance rises or falls on three pillars that repeatedly surface in board and regulator conversations:

1) Transparency (disclosure that is operational, not cosmetic)

Transparency is not a marketing statement; it is evidence that stakeholders can understand where AI is used, what data is being used and with what consent, how frequently data is refreshed, and what limitations are known. KPMG’s Trusted AI checklist recommends designing disclosure mechanisms, documenting limitations, publishing system documentation, and issuing transparency reporting on adoption and remediation.

2) Accountability (clear ownership across the lifecycle)

Accountability requires named responsibility for outcomes, incident triage processes, and auditable oversight. Trusted AI guidance goes further: human-AI interactions should be recorded outside the AI solution for audit and transparency, and Human-in-the-Loop controls should be designed across the lifecycle. This is where governance moves from principle to proof.

3) Responsible data use (data integrity as a governance control)

Predictive AI is only as credible as the data that shapes it. Establishing data catalogues, classifying data, and appointing data owners and stewards are not “data team” tasks; they are governance prerequisites. Trusted AI guidance sets out practical data integrity measures: track and document data origin, history, and transformation; define go/no-go criteria; and continuously monitor outputs after launch.

“Good” Governance is Measurable and Testable

Predictive AI creates a new class of governance metrics that boards should expect to see alongside traditional risk and performance reporting including:

- drift and degradation signals (reliability over time),
- override rates and reasons (human oversight effectiveness),
- transparency artefacts published and updated (disclosure discipline),
- data quality and lineage indicators (integrity and defensibility),
- incident counts, severity, and remediation cycle time (operational resilience).

Industry leading practices emphasizes end-to-end testing prior to deployment and ongoing monitoring, including established metrics, recalibration procedures, and stakeholder reporting plans for transparency.

The Governance Playbook For 2026: Foresight with Confidence

The organizations that win with predictive AI will not be those with the most models. They will be those who can answer, quickly and credibly:

- Which decisions are AI-influenced, and at what autonomy level?
- Who is accountable for outcomes, incidents, and model changes?
- What evidence exists for transparency, explainability, and data integrity?
- How are models monitored, recalibrated, and governed over time?

KPMG’s Trusted AI approach provides a practical backbone for this, built around values-driven, human-centric, and trustworthy principles, underpinned by ten ethical pillars including transparency, accountability, explainability, privacy, security, and data integrity.

For boards, audit committees, and executive teams across the region, the opportunity is to lead the process in setting governance expectations that, allow innovation to scale while protecting trust.


As we move into 2026, the differentiator will be organizations that can evidence not just assert that their AI is governed, explainable, and resilient over time. For organizations seeking to move from intent to implementation, it is critical that you have the appropriate support to guide you from both a technical and a governance perspective; this includes support right from the design stage of the AI governance operating model, embed lifecycle controls, and implement practical tooling and reporting that stands up to scrutiny and supports sustainable value creation.

About the Author

Jeyapriya Partiban is a Partner and the Head of Advisory for KPMG in Bahrain. She is an expert in the realm of Risk Consulting with a specific focus on Governance & Compliance. Having joined KPMG in Bahrain in 2005, Jeyapriya has led and grown the Risk Consulting practice for KPMG; while serving as a trusted advisor for a number of clients with the Public and Private sectors.



**CRESCENT
ENTERPRISES**



Governance architecture, however well designed, only delivers value if it is built on the right assumptions about who AI serves and who it affects. Crescent Enterprises widens the frame considerably, situating internal governance within a much larger question about the societal consequences of AI adoption at scale. Their contribution asks boards to hold two responsibilities simultaneously: accelerating responsibly within their organisations, and remaining accountable for the wider impact of the systems they deploy. It is a perspective that raises the stakes of the governance conversation in ways that purely technical frameworks do not.

Corporate Governance in the Age of Predictive AI: From Insight to Foresight, by Ola Al Haj Hussin

Predictive AI is rapidly becoming a core decision engine as companies' governance expectations shift from hindsight reporting to forward-looking signalling. But prediction without integrated governance can harden bias, blur accountability, and turn probabilistic outputs into unchallengeable decisions.

At Crescent Enterprises, predictive analytics and AI are increasingly embedded across operations, systems, and workflows – from forecasting demand, anticipating supply chain disruptions, spotting cyber signals, strengthening compliance monitoring, and enhancing consistency. That is why governance is not a parallel conversation for us – it is central to how these systems are deployed and trusted.

While predictive systems use statistical modelling and machine learning to identify patterns and forecast likely outcomes, today's prescriptive AI tools go a step further by recommending actions to achieve (or avoid) specific outcomes. With this shift from insight to foresight, governance cannot remain anchored in yesterday's cadence of quarterly dashboards, annual model reviews, and occasional policy updates.

The AI Risk Management Framework of the US Department of Commerce captures this reality well: trustworthy AI requires characteristics like accountability and transparency, explainability, privacy safeguards, and managed bias, supported through ongoing "govern, map, measure, manage" risk management across the lifecycle. In short, when prediction becomes infrastructure, governance must become architecture.

The AI acceleration without inclusion risk

To trust that architecture and infrastructure – within corporate governance frameworks as within wider AI applications – we must recognise the inherent imperfections of seemingly impeccably advanced technologies. After all, whilst we are living through exponential AI advancement we are seeing widening asymmetry. The World Economic Forum has noted that only about 100 of the world's 7,000+ languages are meaningfully represented in today's AI systems, while a third of humanity remains offline completely. If AI is trained on narrow language and data environments, it can systematically misunderstand (or entirely ignore) large segments of customers, workers, and communities. Despite unprecedented global AI investment – more capital was deployed to AI in 2025 alone than to the atomic bomb and moon landing programs combined – less than 1% of total AI investment has been directed towards positive social impacts and most benefits remain concentrated in high-income economies, further raising the stakes of bias, inequality, and power concentration.

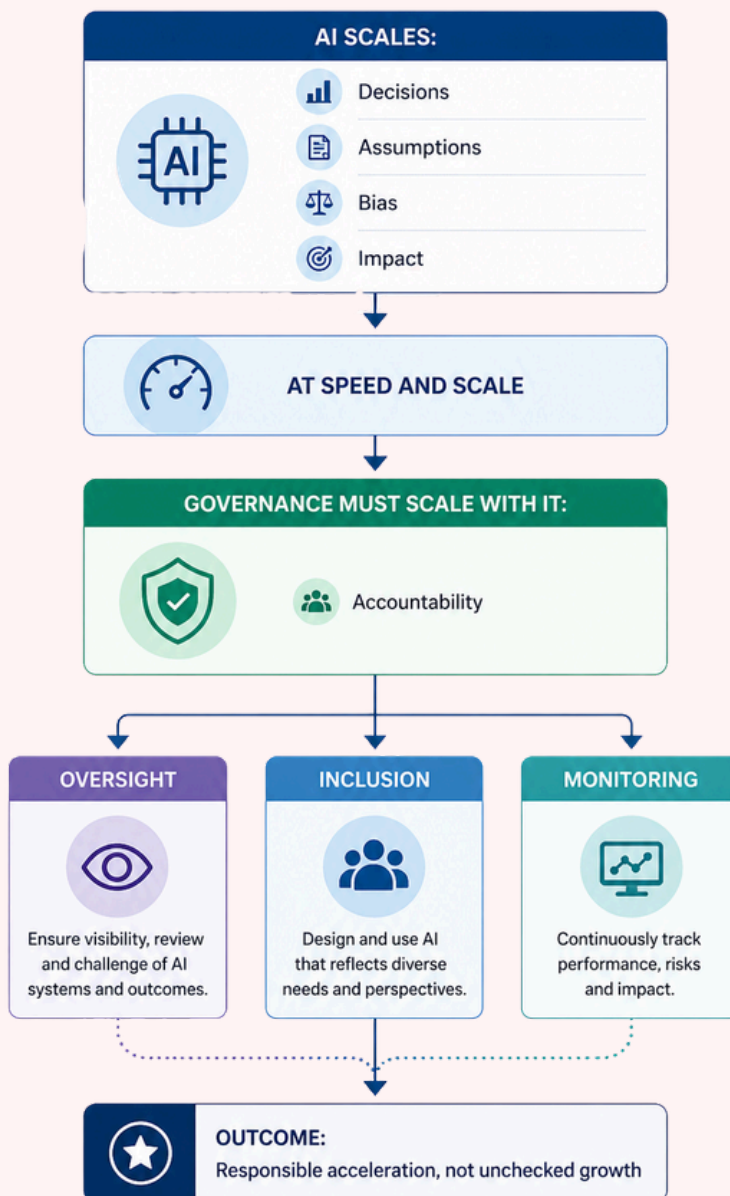
Alongside the risk is opportunity. Crescent Enterprises has been championing the alignment of AI and social innovation to catalyse transformative growth in otherwise isolated sectors. Just as many frontier markets embraced mobile technology, fintech, and blended finance to leapfrog outdated models, the AI acceleration can follow a similar path. Realising it requires connected capital, open-data ecosystems, and cross-sector partnerships that democratize access and scale impact. Crescent has invested across the entire AI operational chain, from data centers and satellite services to interoperability systems and integrated AI stacks. Additionally, Crescent has pioneered efforts to build a global 'AI for social innovation' coalition – bringing together leaders from across government, business and philanthropy – and has launched a world-first AI for Social Innovation prize with MIT Solve that recognises how AI is being applied in novel ways to solve real-world problems. From smarter diagnostics and logistics to hyper-local commerce and education, this practical next wave of AI can deliver scalable benefits across developing sectors. For example, AI-powered early warning systems can flag health and climate risks in vulnerable communities before they become catastrophic, which facilitates new insurance and adaptation mechanisms that, in turn, enable additional financing and innovation.

For corporate governance, this new AI era creates a dual mandate: accelerate responsibly, and govern proactively. Boards cannot treat predictive AI as a simple technology upgrade – it is a transformation of how decisions are made, scaled, and trusted. That requires not only greater technological fluency, but stronger governance systems embedded from the outset.

Frameworks such as UNESCO's global AI ethics standards are clear: risks can compound existing inequalities; systems must be auditable and traceable; and accountability must remain human. But beyond compliance, the real challenge for boards is anticipatory governance – ensuring oversight evolves ahead of capability, not behind it.

In practice, that means boards must do two things at once: lean into AI's transformative potential – and stay ahead of its risks. Because AI does not just scale decisions – it scales the assumptions inside them.

WHAT AI ACTUALLY SCALES



A Board-Level Gap Is Emerging as AI Becomes A Decision Engine

AI adoption is moving faster than oversight maturity. McKinsey’s reports that nearly 90% of organizations use AI in at least one business function, but less than 40% of Fortune 100 companies disclose any form of board oversight of AI. Even more telling, a global survey of directors found that two in three report their boards have “limited to no knowledge or experience” with AI, and nearly one in three say AI does not even appear on board agendas.

This is not a critique of boards, it is a warning of adoption mismatch. Predictive AI changes our governance problem from “Are we compliant?” to “Are we accountable?” It also changes the oversight challenge from periodic review to real-time confidence: do we know where models sit in workflows, what they are optimizing, where they are drifting, and who is responsible when they are wrong? And model drift is not theoretical. Deloitte has documented how point-in-time checkpoints can miss risks that emerge after deployment; and models can degrade as conditions change, requiring continuous monitoring rather than static approvals.

Where Predictive AI Belongs – and Where It Does Not

With operations and investments spanning multiple sectors and geographies, Crescent Enterprises' experience is that predictive AI is most effective where the decision is primarily about pattern detection, prioritization, and early warning – especially in high-volume contexts where humans benefit from better triage. It can be valuable in areas like preventive maintenance, fraud detection, inventory planning, and risk sensing – provided there is documentation, monitoring, and a clear ability to intervene.

But there are domains where predictive AI should not be relied upon as the decider – particularly where rights, livelihoods, or safety are on the line, and where fairness and contestability are non-negotiable. That includes decisions like hiring and firing, credit and access decisions, disciplinary actions, and other high-impact determinations – unless there is meaningful human oversight, explainability appropriate to the context, and a pathway to challenge outcomes.

A model can be all-confident in its forecast without consideration for unfair or harmful outcomes if its training data reflects historical inequity or incomplete representation. An algorithm can advise, but only accountable humans should decide.

Predictive AI can help companies anticipate risks, allocate resources wisely, and unlock productivity – but it will only deliver durable value if governance keeps pace with capability. At Crescent Enterprises this means advancing AI deployment alongside equally rigorous accountability – across investments (including along the full AI value chain), operations (from portfolio companies to internal workflows), and partnerships (such as pioneering AI for social innovation initiatives). Together, these reinforce a model where innovation and inclusion move in parallel, not in tension. Ultimately, success in this AI era is not only technical, it is systemic. Speed without governance creates risk; governance without speed creates irrelevance; the advantage lies in mastering both.

About the Author

Ola Al Haj Hussin has over twenty years of experience in development, philanthropy, and responsible business across the MENA region.

She has worked with AGFUND, UNICEF, and UNDP on inclusive development, institutional partnerships, financial inclusion, and humanitarian response.


At Crescent Enterprises, she oversees Corporate Citizenship, with a focus on sustainability, community investment, programme design, and cross-sector partnerships.

Her work is grounded in responsible governance and inclusive development, including how institutions can use innovation and AI to advance measurable social impact.



CHALHOUB GROUP

SINCE 1955



From the broader picture, the report turns to the specifics of what governance actually looks like inside a large, fast-moving organisation that has made AI a strategic priority. Chalhoub Group's contribution is the most candid account in this section of the choices, trade-offs, and deliberate design decisions that go into building an AI governance model that works in practice. Their experience illustrates a principle that runs through this entire report: effective governance is not imposed on an organisation from outside, it is built from within, shaped by the organisation's structure, culture, and commercial reality.

From Insight to Foresight: Corporate Governance in the Age of Predictive AI By Jeremy Denisty

No one knows what the future holds. But in an era where artificial intelligence is reshaping how businesses operate, compete, and create value, uncertainty cannot be an excuse for inaction. The question facing boards and executive teams today is not whether to engage with AI, but how to govern it in a way that is responsible, value-driven, and fit for the organisation they lead.

There is no universal blueprint for AI governance. How a company approaches it is deeply tied to its strategic positioning: whether it sees itself as a leader or a follower, the competitiveness of its industry, its exposure to digital disruption, and its appetite for risk. A financial institution governing predictive credit models will have fundamentally different priorities from a luxury retailer using AI to personalise the customer experience. Governance frameworks must reflect that reality, not chase a one-size-fits-all standard.

At Chalhoub Group, the largest luxury retailer in the Middle East, we take a pragmatic view. AI must deliver value. We do not innovate for the sake of innovation or to generate headlines. Every initiative must solve a real business problem, and our governance model is designed to ensure exactly that. Let me share a few of the principles that guide us.

AI Governance in Practice: From Idea to Scale



GOVERNANCE PRINCIPLE: PROPORTIONALITY



First, structure matters. Our AI Lab reports directly to Amit Keswani, our Chief Omnichannel and AI Officer, who sits on the Executive Committee. While we maintain very strong collaboration with our IT department – we are even physically co-located — the AI Lab is deliberately not fully embedded within IT. This is an intentional choice. AI is fundamentally a people and behaviour transformation, not just a technology deployment. We need to sit as close as possible to the business, to our brands, and to our customers, while operating as a central service function. This proximity to the business is what allows us to govern AI with commercial relevance, not just technical compliance.

Second, the vast majority of our AI initiatives originate from the business itself. We sit down with brand and department leaders and ask a simple question: “forget about AI, what are your biggest problems?” From there, we assess whether AI is the right solution. This approach makes return on investment more straightforward to demonstrate, drives stronger commitment from the teams who will live with the transformation, and delivers a significantly higher success rate. It also means governance is not an afterthought imposed on reluctant stakeholders, it is embedded from the very first conversation.

Third, we operate with the maximum agility that a large corporate environment can enable. Our AI Lab has the mandate to lead Proofs of Concept: we finance them, lead them, and implement them with external vendors or our internal technology team. The guiding principle is the leanest possible approach at the smallest cost until a Proof of Concept proves its value. Once it does, the brand or department for which the solution was built takes it onto their own P&L. This means only ROI-positive initiatives effectively move into the business, and the AI Lab itself operates against a target success rate. Only when we move to scaling do we invest in onboarding or training additional capabilities.

Across all stages, we run initiatives through a series of internal reviews: Information Security, Data Privacy, Ethics and Compliance, most of them coordinated through our dedicated AI Ethics and Compliance Committee. But the intensity is calibrated: reviews are light-touch for Proofs of Concept, and deep, rigorous, and comprehensive once a solution moves to scale. This graduated approach keeps us agile without compromising on accountability.

An important nuance in how we govern predictive AI is the role we assign to it. When we deploy predictive models, the output serves as an informed baseline, not a final answer. Our teams then adjust the prediction with factors that the model doesn't capture. Predictive AI empowers our people with stronger foundations for decision-making; it does not replace their judgement. As for the boundaries of what we will or will not pursue, our philosophy is not to put barriers around specific use cases. We do not maintain a list of things AI cannot touch. Instead, we calibrate governance intensity to the sensitivity of the use case. When a project involves customer-facing decisions, personal data, or ethical considerations, we multiply checks and balances on all fronts. The principle is not avoidance, but proportionality.

I do not believe there is a single best-practice model for AI governance. You need to find what works for your organisation, given your structure, your strategic priorities, and your needs. Artificial Intelligence is a key pillar of Chalhoub Group's Vision 2033, and our current governance model is delivering results. But because this field moves at extraordinary speed, we are constantly questioning, challenging, and refining how we operate. The key is to not be afraid to take action, but to take it in a nimble, agile way, with a healthy tolerance for failure. Every misstep carries invaluable learning that makes the organisation more future-fit.


In the age of AI, the greatest governance risk is not moving too fast. It is standing still.

About the Author

Jeremy Denisty is the Head of AI at Chalhoub Group, with a background in digital innovation, technology strategy, sales, and marketing. He is the author of *The Virtual Economy*, focused on the evolution of the internet economy and its business applications. He has worked on ideating, strategising, and delivering digital innovation projects over the years in his current role. He drives the Group's AI strategy, leading the design and execution of initiatives that embed AI across functions to deliver speed, efficiency, and business impact.



TAMER



Governance principles, however clearly articulated, ultimately have to be translated into operational systems and technical architecture. Tamer Group's contribution addresses this layer directly, examining how organisations can embed transparency, accountability, and lifecycle oversight into the infrastructure of AI itself. If Chalhoub shows what governance looks like from the leadership and strategy perspective, Tamer shows what it looks like from the ground up, in the systems and processes that make governance real rather than theoretical.

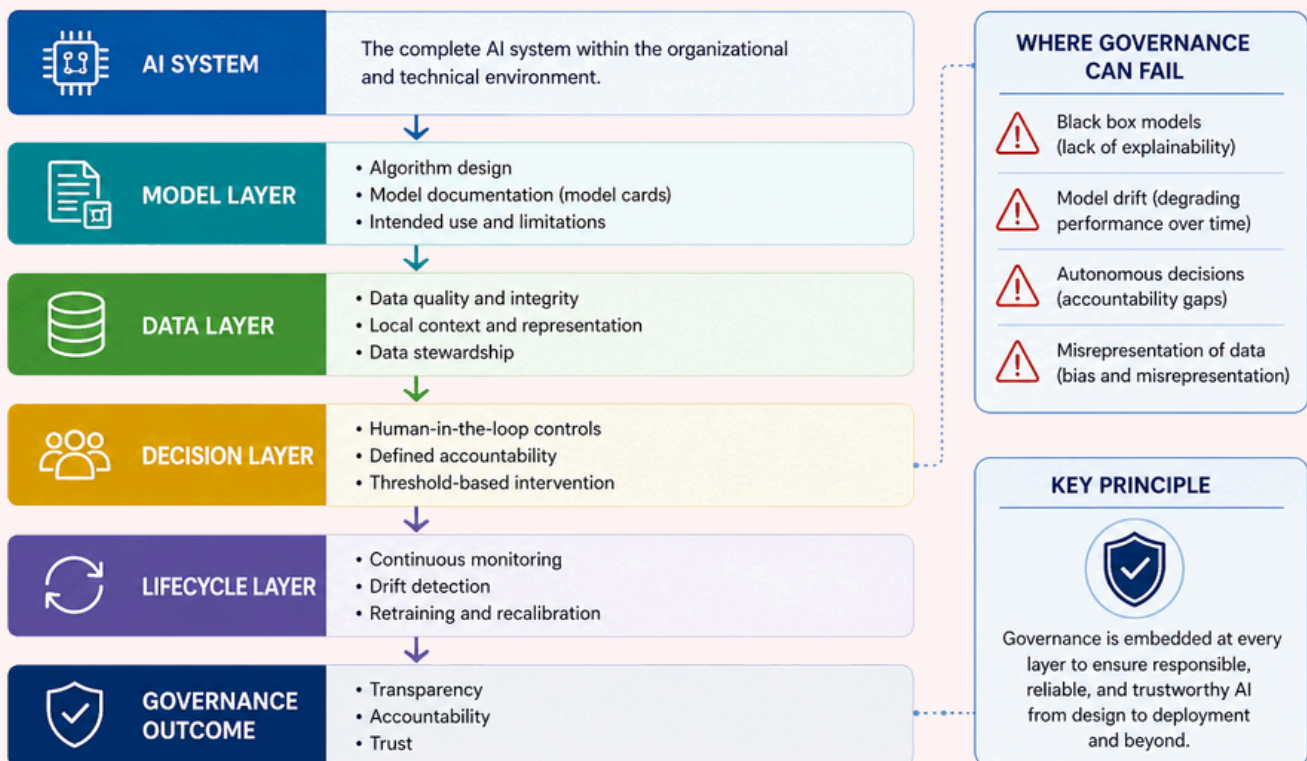
Governing the Algorithm: Building Trust through IT Architecture, by Taha Tkito

Artificial intelligence has moved rapidly from experimentation to operational reality. Across industries, organizations are increasingly using predictive AI to optimize supply chains, forecast demand and support complex operational decisions. In sectors such as healthcare and logistics, these systems are no longer peripheral tools. They are becoming part of the foundational infrastructure of the enterprise.

AI-enabled forecasting can improve demand prediction accuracy by as much as 20%–30%. Yet as algorithmic systems increasingly influence operational decisions, an important governance question emerges: who is responsible for the outcomes?

As regulatory frameworks for artificial intelligence continue to mature globally and across the GCC, organizations cannot rely solely on legislation to define responsible practices. Instead, governance must be embedded directly into how AI systems are designed, deployed and monitored. Governance principles needs to be translated into operational architecture so it does not become not a constraint on innovation but an enabler for AI to scale responsibly.

AI Governance Built into System Architecture



GOVERNANCE IS CONTINUOUS, NOT ONE-TIME.

1. Solving the Opacity Challenge in Predictive AI

One of the central governance challenges is the “black box” nature of advanced AI models. When algorithmic outputs influence high-stakes decisions such as inventory allocation, pricing strategies or resource planning, organizations must ensure that those decisions remain understandable and accountable.

To address this challenge, many organizations are moving toward algorithmic traceability. Documentation frameworks such as model cards or model registries create structured records of how a model was designed, the data it relies on, its intended use cases and its limitations. These mechanisms provide the connective tissue between technical performance and governance oversight allowing auditors, risk committees and leadership teams to verify that an algorithm is functioning as intended.

2. Managing the Lifecycle: The Risk of Model Drift

Governance cannot end at deployment. Predictive models are trained on historical data, but the environments in which organizations operate, particularly in dynamic markets such as GCC, are constantly evolving.

“Model drift” occurs when changes in market conditions, consumer behavior or operational dynamics gradually reduce a model’s accuracy or introduce unintended bias. Without appropriate monitoring, systems that initially perform well may begin to produce unreliable outcomes over time. Effective AI governance therefore requires lifecycle oversight. Monitoring systems must be in place to track model performance continuously and trigger reassessment when deviations occur. By treating algorithms as living assets rather than static tools, organizations can detect performance degradation early and implement corrective actions such as retraining, recalibration or model retirement before operational risk escalates.

3. Agentic AI and the Accountability Gap

The next frontier in enterprise AI is the deployment of agentic systems. These are algorithms that not only generate recommendations but also execute actions autonomously within operational workflows. While these capabilities offer significant efficiency gains, they also introduce new governance challenges.

Agentic systems can create an accountability gap when responsibility for outcomes becomes distributed across data providers, developers and operational users. Without clear ownership, organizations may struggle to determine who is accountable when automated decisions produce unintended consequences.

To mitigate this risk, leading organizations are embedding human-in-the-loop mechanisms into their AI architecture. When automated decisions exceed defined financial, operational or safety thresholds, human oversight remains part of the decision process. In addition, many organizations now assign an executive sponsor to each high-impact algorithmic system, ensuring that responsibility for performance and risk management remains clearly defined.

4. Data Stewardship and Local Context

Algorithms ultimately reflect the data on which they are trained. In diverse markets such as those across the GCC, datasets that lack linguistic or demographic representation can produce unintended bias or inconsistent outcomes.

Responsible AI deployment therefore requires strong data stewardship such as auditing datasets, identifying proxy variables that may unintentionally encode bias and ensuring that training data reflects the diversity of the environments in which systems operate. Ensuring that AI systems are locally grounded and contextually aware is essential for maintaining credibility and trust.


Conclusion: Governance as an Enabling Architecture

The governance of predictive and agentic AI is ultimately about building the trust required for innovation to scale. As artificial intelligence continues to reshape industries and is projected to contribute hundreds of billions of dollars to Middle Eastern economies over the coming decade, organizations that embed transparency, accountability and lifecycle oversight into their AI architecture will be best positioned to capture that value responsibly.

About the Author

Taha Tkito is a seasoned IT executive with over 25 years of experience leading technology organizations across multinational companies in the FMCG, Retail, Healthcare, Logistics and Technology sectors throughout EMEA.

His expertise includes spearheading large-scale digital transformations, building high-performing customer-centric teams and enhancing organisational digital capabilities and operational resilience.



CONSOLIDATED CONTRACTORS COMPANY (CCC)

The governance challenges shift again when AI operates not in a commercial or consumer context, but in a project-driven environment where the consequences of a poor decision can be immediate, physical, and irreversible. Consolidated Contractors Company's contribution brings that reality into focus, describing an approach to predictive AI governance that is deliberately conservative, tightly bounded, and anchored in clear human accountability at every stage. It is a reminder that the appropriate governance response is always calibrated to context, and that in high-risk environments, the cost of getting it wrong is not measured in reputation alone.

Governance of Predictive AI in Decision-Making at CCC, by Rami Ziadeh

At Consolidated Contractors Company (CCC), predictive artificial intelligence (AI) is viewed as a means of strengthening governance and improving early visibility of emerging risks. It is not intended to replace professional judgment or shift accountability. In a project-driven and high-risk operating environment, CCC's priority is to ensure that AI-supported decision-making reinforces the company's long-standing principles of integrity, transparency, hierarchy, and responsibility.

Where Predictive AI Influences Decisions at CCC

Within CCC, predictive analytics and AI-enabled tools are primarily deployed as decision-support mechanisms in areas already governed by established oversight and internal control processes.

These include:

Enterprise and project risk management

Predictive indicators support the earlier identification of emerging safety, delivery, or commercial risks within project and enterprise risk registers. Insights generated through analytics contribute to periodic risk reviews and management discussions.

Health, Safety, Environment & Security (HSE&S)

Trend analysis across inspections, observations, and incident data helps management identify patterns and prioritise preventive interventions in support of CCC's objective of zero incidents.

ESG and sustainability performance

Predictive insights assist in identifying potential underperformance against environmental or social targets before issues become material, allowing earlier corrective action.

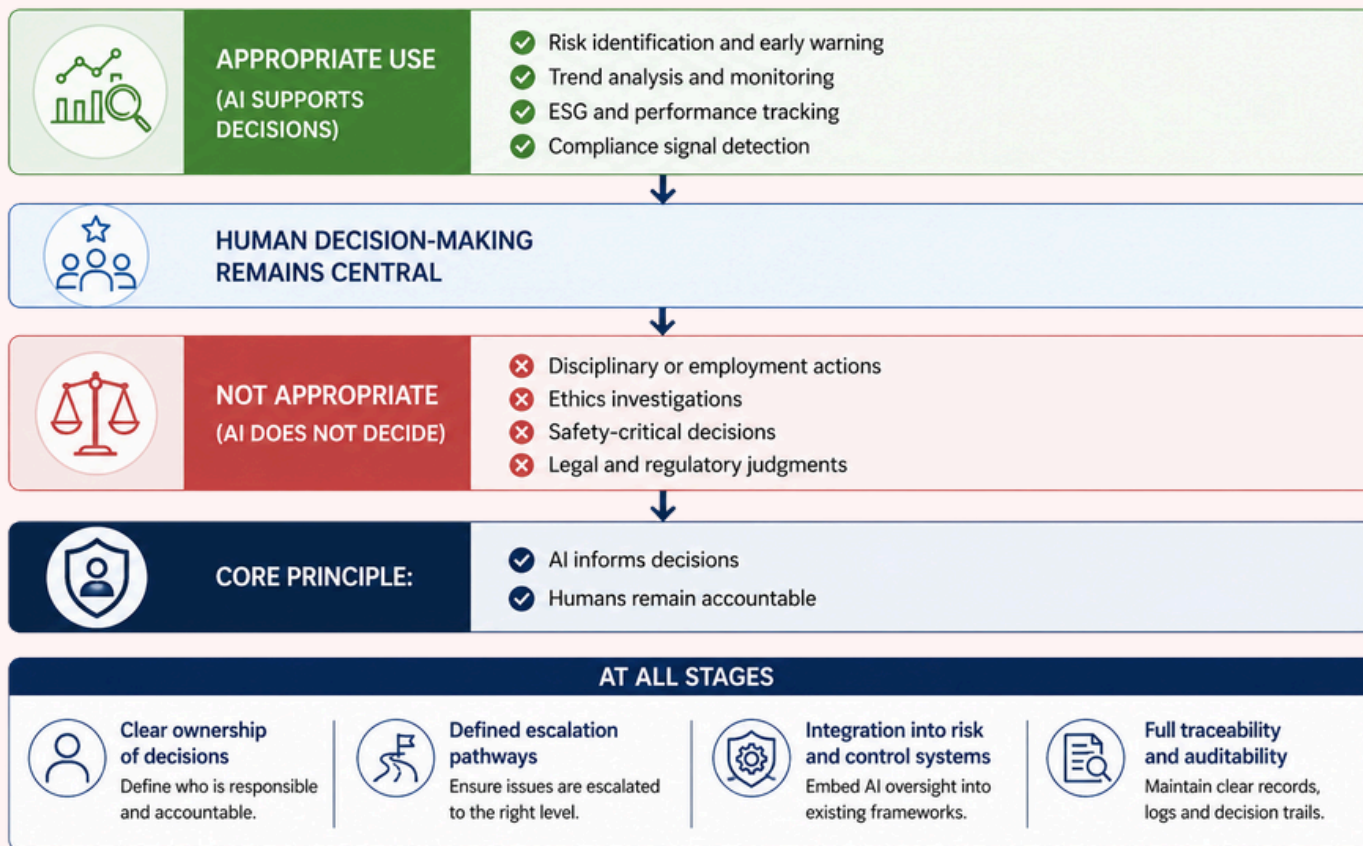
Ethics, compliance, and assurance

Analytical tools may support the detection of anomalies or patterns that warrant management attention. These tools complement, but do not replace, established investigation, escalation, and assurance procedures.

In all cases, predictive AI is used to inform managerial judgment rather than replace it.

Decisions Not Considered Appropriate for Reliance on Predictive AI

Defining Boundaries for Predictive AI Use



CCC does not rely on predictive AI for decisions involving significant human impact, ethical judgment, or heightened reputational or legal sensitivity.

These include, but are not limited to:

- Disciplinary or employment actions
- Ethics investigations
- Safety-critical operational decisions on project sites
- Matters involving legal or regulatory exposure

This approach reflects CCC's commitment to accountability, fairness, and transparency. Decision-making must remain clearly attributable to responsible individuals, with a defensible rationale. Where an AI output cannot meet appropriate standards of explainability and justification, it is treated strictly as advisory input.

CONSOLIDATED CONTRACTORS COMPANY (CCC)

Governance Oversight and Escalation

Predictive AI at CCC is governed through the company's existing executive and management governance structures, rather than through standalone technology forums. Oversight is embedded within established risk management, ethics, compliance, and sustainability governance arrangements.

Management is expected to maintain clear visibility over:

- where predictive AI is used,
- which decisions it may influence,
- what controls and approvals apply, and
- how exceptions, failures, or unintended outcomes are escalated.

This approach ensures that AI-supported activities remain fully integrated within CCC's established reporting, review, and escalation processes.

Accountability for Outcomes

Accountability for decisions influenced by predictive AI remains with business and operational leadership. While technology functions are responsible for system integrity and performance, and risk or compliance functions provide challenge and assurance, ultimate responsibility for decisions and their consequences rests with accountable management.

This principle reflects CCC's long-standing emphasis on clear hierarchy and managerial responsibility.

Risk Appetite and Boundaries

CCC adopts a measured and conservative risk appetite for the use of predictive AI.

Acceptable use cases are those that strengthen early warning capabilities, monitoring, and prioritisation within well-controlled processes. Conversely, unacceptable use cases include those that:

- automate high-impact or sensitive decisions,
- create ambiguity over accountability, or
- expose the organisation to material reputational or regulatory risk.

Clear boundaries and governance controls are expected to be defined before any material deployment of predictive AI capabilities.

Integration into Enterprise Risk Management and Internal Controls

Predictive AI operates within CCC's Enterprise Risk Management (ERM) and internal control framework, rather than alongside it.

This includes:

- documented use cases,
- defined approval and oversight requirements,
- model validation and testing,
- ongoing performance monitoring, and
- periodic governance review.

Where predictive AI materially influences risk exposure or control effectiveness, this influence is expected to be reflected in management reporting and risk review processes.

Explainability, Auditability, and Traceability

Explainability requirements are applied proportionately to the potential impact of the decision context. High-impact areas—including safety, ethics, workforce matters, and external reporting—require clear rationale, documentation, and human approval.

AI-influenced decisions must remain auditable and traceable. CCC's established discipline in evidence management provides the foundation for this approach. Relevant records may include:

- model versions,
- data sources,
- assumptions and thresholds,
- approval documentation, and
- decision logs.

Maintaining this documentation ensures that outcomes can be reviewed, challenged, and defended where necessary.

Value Delivered and Key Lesson

Predictive approaches have demonstrated value by improving foresight and accelerating escalation, enabling management to intervene earlier in emerging risk or compliance trends.

A key lesson is that predictive tools are only as effective as the governance, data quality, and contextual judgment surrounding them. Without strong controls and disciplined oversight, predictive outputs can create a false sense of certainty rather than meaningful clarity.

Conclusion

CCC supports the responsible use of predictive AI where it strengthens oversight, enhances risk visibility, and reinforces disciplined decision-making.

The adoption of predictive AI does not alter the fundamental principles of governance at CCC. Leadership remains accountable, decisions remain explainable, and the protection of people, reputation, and corporate values remains paramount.


About the Author

Rami Ziadeh is a Governance, Compliance & Special Projects Manager with 20+ years of international experience supporting large-scale infrastructure and energy projects. Proven expertise in regulatory compliance, strategic reporting, and cross-functional coordination across Europe, Middle East, Africa, and Australia.



EY

**Building a better
working world**



The final contribution steps back from organisational experience to ask what all of this means for governance as a discipline. EY's reflection reframes the journey traced across this section, from architecture to practice to high-stakes application, as part of a deeper transformation in how accountability itself must be understood in an AI-enabled world. It is a fitting close to the section, and a bridge to the recommendations that follow.

The Shifting Role of Predictive AI in Governance, by Hisham Dally

As predictive Artificial Intelligence (AI) becomes increasingly embedded in core business processes, organizations are transitioning from hindsight-driven analysis to forward-looking, data-enabled decision-making. Across financial advisory, Internal Controls over Financial Reporting (ICFR), and board governance engagements, this shift is becoming increasingly evident. AI is no longer peripheral; it is actively shaping financial judgments, risk assessments and strategic direction. This evolution requires a corresponding transformation in how governance frameworks are designed and applied.

Predictive AI introduces a new layer of complexity to governance because it challenges the foundational principles of transparency, accountability and control. Traditional governance frameworks, particularly those underpinning ICFR, are built on deterministic processes, auditable trails and clearly attributable decisions. By contrast, AI models often operate probabilistically, evolve over time and depend on large, dynamic datasets. This creates inherent tension between innovation and control.

Implications for financial reporting

From an ICFR perspective, the integration of predictive AI into financial processes raises critical questions. How do we validate the integrity of model outputs that inform financial estimates or provisions? How do we provide evidence of controls over models that continuously learn and update? Governance frameworks must expand to include model validation controls, data integrity checks, and clear documentation of assumptions and limitations. Without this, organizations risk undermining the reliability of financial reporting — an area where trust is paramount.

Transparency is a key challenge. Boards and audit committees are increasingly confronted with decisions influenced by AI, yet may not fully understand how those outputs are generated. In practice, this creates a “black box” risk, where decision-makers rely on insights they cannot fully interrogate. Addressing this requires embedding explainability into governance processes. It does not necessarily simplify complex models, but it facilitates the communication of their logic, inputs and limitations in a way that supports informed oversight. When it comes to board support, there is a growing demand for translating technical complexity into actionable governance insight.

Redefining Accountability in an AI Enabled Governance Model

Accountability must also be redefined. In traditional governance structures, accountability is clearly delineated across management, control functions and the board. With AI, accountability can become diffused across data owners, model developers and business users. Organizations must establish clear ownership of AI-driven decisions, including defined roles for model governance, risk oversight and escalation. Importantly, human judgment must remain central. AI should augment, not replace, decision-making, particularly in high-stakes financial and strategic contexts.

In practice, this shift is prompting boards and audit committees to ask a new set of questions. Which decisions are being influenced by predictive AI, particularly in financial and risk-related areas? How are models validated, monitored and periodically recalibrated? Who ultimately owns the outcomes of AI-driven decisions, and how are exceptions escalated? Critically, how do organizations confirm that increased reliance on AI does not dilute accountability, but rather strengthens informed decision-making?

Addressing these questions ultimately depends on strong data governance. Predictive AI is fundamentally dependent on the quality, completeness and integrity of data. Weaknesses in data governance can propagate through models and lead to flawed outputs, with direct implications for financial reporting, compliance and reputation. Strengthening data governance, including lineage, access controls and validation, must therefore be a priority. From a governance standpoint, this is not just a technical issue, but a board-level concern given its enterprise-wide impact.

At the board level, the implications are significant. Boards are increasingly expected to oversee not only financial and operational risks, but also technology-driven risks and opportunities. This requires enhancing board capabilities, either through upskilling or by bringing in specialized knowledge, to facilitate effective oversight of AI. In many cases, audit and risk committees are expanding their mandates to include AI governance, reflecting its intersection with financial reporting, controls and enterprise risk management.

Transitioning to Continuous, Real-Time Oversight

Another critical shift is the move from periodic to continuous oversight. AI systems operate in real time and evolve dynamically. This means that traditional governance cycles are no longer sufficient. Organizations should implement continuous monitoring mechanisms, including model performance tracking, exception reporting and periodic revalidation. This aligns closely with the evolution we are seeing in ICFR, where there is a growing emphasis on ongoing controls monitoring rather than point-in-time assessments.

Finally, governance in the age of predictive AI must be anchored in responsibility and ethics. Beyond compliance, organizations must consider the broader implications of AI-driven decisions, particularly in areas such as bias, fairness and unintended consequences. Embedding responsible AI principles into governance frameworks is essential to maintaining stakeholder trust. This includes clear policies, training and a culture that encourages critical challenge and accountability.

The Shift in Governance

TRADITIONAL GOVERNANCE			AI-ENABLED GOVERNANCE	
	Hindsight-based	➔		Foresight-driven
	Periodic oversight	➔		Continuous monitoring
	Deterministic processes	➔		Probabilistic models
	Clear decision ownership	➔		Distributed accountability
	Static controls	➔		Dynamic, evolving systems

Governance must evolve to ensure

	Key Requirement	Establish fundamental governance standards and expectations.		Human judgement remains central	Keep humans in control of critical decisions and ethical judgement.
	Transparency and explainability	Ensure AI decisions can be understood, explained, and communicated.		Distributed accountability	Share accountability across teams, systems and the value chain.
	Clear ownership of outcomes	Define who is accountable for decisions and their impacts.		Dynamic, evolving systems	Adapt governance as AI, risks and contexts continuously evolve.
	Strong data governance	Ensure data quality, privacy, security and appropriate use.		Continuous monitoring	Move from periodic reviews to ongoing, continuous oversight.
	Real-time oversight mechanisms	Enable real-time visibility, alerts and issue detection.		Clear decision ownership	Ensure accountability is clear, traceable and audit-ready.

Framing The Path Forward

The transition from insight to foresight represents a fundamental shift — not only in how decisions are made, but also in how they are governed. Organizations that succeed will be those that integrate AI into their governance frameworks with rigor and discipline while preserving the core principles of transparency, accountability and control. From an ICFR and board governance perspective, this is an opportunity to redefine governance for a more complex, data-driven future — so that innovation is matched with trust and resilience.

About the Author

Hisham Dally is the EY MENA FAAS ICFR and Governance Solutions Leader. He leads a diversified pool of talented professionals across the MENA region to deliver a wide range of services in the areas of Internal Controls over Financial Reporting (ICFR), governance and board support. With over 16 years of experience, Hisham has led transformational projects and championed strategic initiatives across multiple sectors and industries. Hisham works closely with board members and C-suite executives to strengthen governance structures, elevate financial integrity, and embed robust control environments that drive accountability towards sustainable value creation. Prior to his current role, Hisham established and led multiple governance, finance control, risk and compliance functions.

INDUSTRY PERSPECTIVES

Across these six perspectives, several things emerge with consistency. Every contributor, regardless of sector or organisational type, draws a firm line between AI as a tool that informs decisions and AI as a mechanism that makes them. Human judgment, they agree, is non-negotiable, not as a symbolic safeguard but as a genuine accountability requirement. Every contributor also recognises that governance cannot be a one-time exercise. AI systems evolve, data changes, and models drift. Oversight must be continuous, not periodic.

Where the contributions differ is in how that governance is built and where it sits within the organisation. There is no single model that works across contexts, and none of the contributors claim otherwise. What they share is the conviction that the work of building governance cannot wait for regulation to arrive and define it for them.

That conviction is, perhaps, the most important thing this section offers. The organisations represented here are not waiting. They are building, learning, and in some cases leading. For boards across the region still weighing whether and how to engage with this agenda, that is both a prompt and a precedent.

What Organisations Governing Predictive AI Are Learning

- The line between AI that advises and AI that decides can blur quickly. Defining it deliberately, and revisiting it regularly, is one of the most important things a board can do.
- Governance that is designed around compliance will always lag behind capability. The organisations building durable frameworks are designing for accountability, not just regulatory readiness.
- Human judgment is not a fallback for when AI fails. It is a structural requirement, built into the decision-making process from the outset.
- Accountability does not transfer to a model. When an AI-influenced decision goes wrong, responsibility still sits with the people and institutions that deployed it.
- Governance intensity should be proportionate to decision impact. Not every use case requires the same level of scrutiny, but every high-stakes decision requires a named human owner.
- Models degrade. Market conditions change, data shifts, and systems trained on historical patterns can produce unreliable outputs over time. Continuous monitoring is not optional, it is the governance baseline.
- Data governance is AI governance. Poor data quality does not stay in the data team. It travels through the model and into the decision.
- Vendors and third-party tools do not sit outside the governance framework. If your organisation deploys it, your organisation is accountable for it.
- Speed and governance are not in tension. The organisations moving fastest with AI are, in many cases, the ones that invested earliest in the structures that allow them to move with confidence.
- AI adoption without inclusion creates risk at scale. Systems trained on narrow data environments can systematically misrepresent the communities and markets they are meant to serve.

**FROM
PRINCIPLES TO
PRACTICE:
GOVERNING
PREDICTIVE AI**

GOVERNING PREDICTIVE AI

1. AI Oversight Structure

AI oversight should be formally embedded within existing board and executive structures. This includes clear reporting lines, defined accountability across the AI lifecycle, and established escalation protocols. Where AI influences strategic or high-impact decisions, oversight must sit at the board level, not solely within technical teams.

2. AI Decision Mapping

Organisations must maintain visibility over where AI is influencing decisions. This includes mapping use cases, classifying levels of autonomy, and identifying high-impact decisions. Governance begins with visibility. Boards cannot oversee what is not clearly understood or documented.

3. Human Judgement Control Points

Human judgement must remain embedded in high-stakes decisions. Clear thresholds should define where human intervention is required, alongside documented override mechanisms and escalation processes. Accountability cannot be delegated to systems.

4. Governance Metrics

Effective AI governance must be measurable. Boards should receive regular reporting on model performance, override patterns, incidents, and data quality. Oversight must be supported by evidence, not assumption.

Board Recommendations

1. Establish Visibility Over AI Use

Boards must ensure there is a clear, organisation-wide view of where predictive AI is being used and which decisions it influences.

- Maintain a register of AI use cases
- Classify use by level of decision impact (advisory vs automated)

2. Define Accountability for AI-Influenced Decisions

Every AI-influenced decision must have a clearly identified human owner.

- Assign accountability at the business level, not just technical teams
- Define escalation and incident response protocols

3. Set Clear Boundaries for Human Oversight

Boards should define where human judgment is mandatory and cannot be overridden.

- Identify high-risk decisions (financial, legal, ethical, reputational)
- Establish human-in-the-loop requirements

4. Integrate AI into Existing Governance Structures

AI governance should not sit separately, but be embedded into existing board and committee oversight.

- Include AI in risk, audit, and compliance discussions
- Ensure regular reporting on AI-related risks and performance

GOVERNING PREDICTIVE AI

5. Strengthen Data Governance as a Foundation

The reliability of AI outputs depends on the quality and governance of underlying data.

- Define data ownership and stewardship
- Monitor data quality, provenance, and limitations

6. Govern Third-Party AI and Vendor Risk

Boards must ensure that externally sourced AI tools meet the organisation's governance standards.

- Require transparency on how models operate and use data
- Extend governance expectations to vendors and partners

7. Build AI Literacy at Board and Leadership Level

Effective oversight requires a baseline understanding of how AI systems function and where risks arise.

- Provide targeted briefings for board members
- Focus on decision impact, not technical detail

Conclusion

Predictive AI is no longer a future consideration. It is already embedded in how organisations make decisions, allocate resources, and manage risk. The question for boards is no longer whether to engage with AI, but whether they are exercising corporate governance with the level of rigour these systems now require.

Across this report, one point is clear. The challenge is not technical, it is one of corporate governance. It is about ensuring that accountability remains visible, that human judgement remains embedded, and that oversight keeps pace with systems that are continuously evolving.

There is no single model for governing predictive AI. What works will differ across sectors, organisational structures, and risk profiles. But the underlying corporate governance principles are consistent. Clarity on where AI is used. Defined ownership of outcomes. Proportionate oversight based on decision impact. And a recognition that corporate governance must be continuous, not periodic.

The organisations represented in this report are not waiting for regulation to define these approaches. They are building corporate governance frameworks in real time, adapting as they learn, and in doing so, setting the direction for what responsible AI adoption looks like in practice.

For boards across the region, the implication is straightforward. Corporate governance cannot follow AI adoption. It must evolve alongside it.

Those that act early will not only reduce risk, but build the institutional confidence to use AI effectively. Those that delay may find themselves accountable for systems they do not fully understand, and decisions they cannot fully explain.

In the age of predictive AI, corporate governance is not a constraint on progress. It is what makes progress sustainable.


GOVERNING PREDICTIVE AI

This is fundamentally a corporate governance issue. It reinforces a principle boards have long understood but must now actively revisit:

Corporate governance is not a design exercise that can be set once and left to operate on autopilot. It requires continuous vigilance, active oversight, and the willingness to adapt as conditions change.

”

PARTNER CONTRIBUTORS



Hisham Dally

MENA FAAS ICFR and Governance Solutions Leader
EY



Shape the future
with confidence



Jeremy Denisty

Head of AI
Chalhoub Group



CHALHOUB GROUP
SINCE 1955



Jeyapriya Partiban

Partner & Head of Advisory
KPMG Bahrain



Ola Al Haj Hussin

Corporate Citizenship Manager
Crescent Enterprises



Rami Ziadeh

Governance, Compliance & Special Projects Manager
Consolidated Contractors Company



Taha Tkito

Chief Information Officer
Tamer Group



PEARL INITIATIVE CONTRIBUTORS

Ishanya Nadkarni

Associate - Governance in Philanthropy Programme
Pearl Initiative



Nada Abdulhadi

Communications Manager
Pearl Initiative



Noor Zuberi

Senior Associate - Governance in Philanthropy Programme
Pearl Initiative



Vineetha Mathew

Senior Advisor - Programmes
Pearl Initiative



Corporate Governance Insight Series

Edition 1

CONTRIBUTORS



Shape the future
with confidence



TAMER



CHALHOUB GROUP
SINCE 1955

