

Corporate AI Governance

Stewardship, Strategy,
and Accountability
in the Age of AI

Release 1.0
May 2026



The Corporate AI Governance Code

Stewardship, Strategy, and Accountability in the Age of AI

Release 1.0

May 2026

This document sets out 'The Corporate AI Governance Code'. The Code is new and, at the date of publication, has no formal regulatory status; it is proposed as best practice for adoption by companies, and for consideration by regulators, governance bodies, and institutional investors as a framework for the oversight and standardised comparison of AI governance at board level.

Preface

Artificial intelligence represents the most consequential general-purpose technology revolution deployed in business since the advent of the internet. It reshapes how organisations create value, make decisions, manage risk, and relate to their workforces and customers. It is not a discrete technology programme to be delegated to an individual function alone. It is a fundamental redesign of the enterprise.

Boards that govern AI well can create a durable competitive advantage, protect their stakeholders, and build the trust on which sustainable success depends.

This Code is proposed in recognition of three developments that now make specific AI governance guidance necessary:

- AI systems are increasingly making or materially influencing decisions that affect customers, employees, counterparties, and the public.
- The risks of ungoverned AI, including bias, opacity, dependency, and systemic failure, are now material in the financial and reputational sense.
- Existing corporate governance frameworks, while applicable, do not provide sufficient specificity to guide boards in discharging their AI-related duties.

This Code does not displace other codes or guidance, or the requirements of existing regulators. It supplements them, providing the AI-specific framework that boards, audit committees, remuneration committees, and their advisers require.

Table Of Contents

Scope and Application	1
Section A – Leadership and Purpose.....	2
Principle A.1 – Board Ownership of AI Strategy	2
Principle A.2 – Long-Term Agency and Technological Dependency	2
Section B – Division of Responsibilities.....	3
Principle B.1 – Board Accountability for AI Governance	3
Principle B.2 – Board Competence	3
Section C – Decision-Making and Control.....	4
Principle C.1 – AI Decision Architecture.....	4
Principle C.2 – AI Workflow Integration	4
Section D – Risk, Assurance and Internal Control.....	5
Principle D.1 – AI as Enterprise Risk.....	5
Principle D.2 – Internal Control and Assurance	5
Principle D.3 – Data Governance	6
Section E – People, Culture and Remuneration.....	7
Principle E.1 – Workforce Impact and Transition	7
Principle E.2 – Remuneration and Incentives	7
Principle E.3 – Culture and the Workforce Contract	7
Section F – Transparency and Accountability	8
Principle F.1 – AI Governance Disclosure.....	8
Principle F.2 – Accountability and Redress	8
Closing Statement.....	9

Scope and Application

This Code applies to any company, but it is particularly relevant to large private companies. The Code applies on a proportionate basis. Companies should apply each Provision in a manner appropriate to the scale, complexity, and materiality of their AI activity, the size and resources of the organisation, and the sectors in which they operate.

The board's judgement on proportionality should itself be documented and, where material, disclosed. Proportionality is not a means of avoiding the discipline the Code requires; it is a means of applying that discipline in a form appropriate to the specific company.

The Code is intended to operate on a comply-or-explain basis. Companies are expected to comply with all provisions of the Code or explain in their annual report how alternative arrangements achieve equivalent outcomes and why those arrangements are appropriate for the company's circumstances.

Explanations should be specific and substantive. Generic explanations that do not enable shareholders to evaluate the board's approach will not satisfy this requirement.

Section A – Leadership and Purpose

A well-run company is one where the board sets clear direction, takes responsibility for outcomes, and ensures that AI serves the long-term interests of the company and its stakeholders. AI amplifies both the ambition and the risk of any organisation. The board must own both.

Principle A.1 – Board Ownership of AI Strategy

The board should ensure that AI is treated as a matter of strategic business importance and is not solely managed as a technology initiative.

Provisions

A.1.1 Strategic Integration. AI should be treated as part of corporate direction, not as an optional innovation track. The board should decide the role AI plays in long-term value creation and ensure that the company's posture on AI is aligned with purpose, strategy, culture, stakeholder expectations, and the economics of the business.

A.1.2 Executive Ownership. A named executive should carry enterprise-wide accountability for AI strategy and governance. This should not displace other leaders' responsibilities, but it should ensure there is a single point of executive answerability to the board.

A.1.3 AI as Business Redesign. The board should direct AI investment toward redesigning material processes and decisions, not toward a permanent state of pilots. AI should alter how the business works, not simply decorate the existing operating model.

Principle A.2 – Long-Term Agency and Technological Dependency

The board should actively manage the company's dependency on AI systems and external AI providers so as to preserve strategic control and operational resilience.

Provisions

A.2.1 Dependency Assessment. The company should know where its AI capability depends on external models, platforms, data sources, infrastructure, and specialist talent, and which of those dependencies are material.

A.2.2 Resilience and Control. Where the company depends materially on third-party AI, the board should ensure there are contractual, technical, and operational protections against loss of access, deterioration in service, or strategic disadvantage.

A.2.3 Strategic Optionality. The board should protect the company's ability to compete, adapt, and retain agency as AI markets evolve. That includes avoiding arrangements that quietly transfer proprietary know-how, data value, or workflow control to external providers.

Section B – Division of Responsibilities

There should be a clear division of responsibilities across the board and executive management for AI governance, with no individual exercising unfettered authority over AI strategy, deployment, or oversight.

Principle B.1 – Board Accountability for AI Governance

The board should establish clear lines of accountability for AI across the organisation and satisfy itself that these are operating effectively.

Provisions

B.1.1 Governance Structure. The board should decide how AI governance is distributed across the board and its committees, ensuring the arrangement is clear, coherent, and appropriate to the company's circumstances.

B.1.2 Management Accountability. Management responsibilities for AI strategy, delivery, risk, ethics, data, and oversight should be explicit and visible to the board.

B.1.3 Group and Cross-Border AI Governance. The board should ensure that AI governance standards operate coherently across the group, including subsidiaries, shared services, outsourced arrangements, and cross-border operations, and that local variation is controlled, justified, and visible.

Principle B.2 – Board Competence

The board should ensure it possesses sufficient knowledge and understanding of AI to discharge its governance responsibilities effectively.

Provisions

B.2.1 Expertise and Advice. The board should assess whether it has enough collective capability to oversee AI and should address gaps through development, appointments, or formal advice.

B.2.2 Board Briefings. The board should receive regular, structured information on the company's AI activity, external developments, regulatory change, and major incidents or risks.

B.2.3 Independent Scrutiny. Non-executives should be able to challenge AI-related proposals and should have access to independent expertise when needed.

B.2.4 Board and Director Use of AI. The board should govern its own use of AI and ensure that directors, board committees, and the board support process use AI tools in a manner consistent with confidentiality, privilege, recordkeeping, and sound judgement.

Section C – Decision-Making and Control

AI systems are increasingly embedded in the decisions that shape a company's strategy, its treatment of customers, and its relationship with employees. The board must ensure that decision rights are clearly defined, that human accountability is preserved, and that appropriate controls govern AI-influenced outcomes.

Principle C.1 – AI Decision Architecture

The board should ensure the company has a clear and documented framework governing how AI informs, supports, or makes decisions, with defined human accountability for all material outcomes.

Provisions

C.1.1 Decision Classification. Management should maintain a register of material decisions and classify the role of AI in each, distinguishing advisory, automated, and hybrid uses.

C.1.2 Human Accountability. Every material AI-assisted or AI-automated decision should have a named human accountable for outcomes.

C.1.3 Override and Escalation. Material AI systems should have functioning mechanisms for human override, escalation, and intervention, and those mechanisms should be tested.

C.1.4 High-Stakes Decisions. Where AI informs decisions with significant consequences for individuals, the company should provide documented human review processes and clear communication to affected parties.

C.1.5 Autonomous and Agentic Operations. Where AI systems can initiate, sequence, or execute actions with limited real-time human direction, the company should define and control the authority, boundaries, supervision, and intervention arrangements governing those operations.

Principle C.2 – AI Workflow Integration

The board should ensure that AI is embedded in the company's material business processes and that its value is measured against operational outcomes, not solely against capability metrics.

Provisions

C.2.1 Operational Deployment. Boards should expect AI investments to be integrated into real workflows with performance metrics that link AI use to business outcomes.

C.2.2 Portfolio Discipline. AI initiatives should be managed as a portfolio with clear priorities, deployment pathways, and decisions to scale, redirect, or stop work.

Section D – Risk, Assurance and Internal Control

AI introduces a distinct category of risk that does not map cleanly onto traditional enterprise risk frameworks. The board has primary responsibility for ensuring that AI risks are identified, assessed, and controlled appropriately throughout the technology lifecycle.

Principle D.1 – AI as Enterprise Risk

The board should ensure that AI risks are integrated within the company's enterprise risk management framework and that material AI risks are reported to the board on a regular basis.

Provisions

D.1.1 Risk Integration. AI risks should be assessed within the enterprise risk framework, with a taxonomy capable of capturing AI-specific characteristics.

D.1.2 Material Risk Disclosure. Where AI risks are material, they should appear in annual report risk disclosures and not remain hidden in internal technical reporting.

D.1.3 Risk Appetite. The board should articulate the kinds of AI use it is willing to tolerate and the degree of error, opacity, bias, and dependency acceptable in each context.

Principle D.2 – Internal Control and Assurance

The company should establish and maintain robust systems of internal control over AI systems and should provide the board with credible assurance that these controls are operating effectively.

Provisions

D.2.1 Lifecycle Controls. AI systems should be controlled across development, procurement, deployment, monitoring, change, and retirement.

D.2.2 Intervention Capabilities. The company should be able to suspend, override, or modify material AI systems when failures, unintended outputs, or changed conditions arise.

D.2.3 Incident Management. AI-related incidents should be identified, escalated, investigated, and remediated through a defined process, with significant matters reported promptly to the board.

D.2.4 Third-Party AI. Before using third-party AI, the company should perform due diligence and ensure ongoing monitoring obligations are built into contractual and operational arrangements.

D.2.5 AI Security, Red-Teaming and Abuse Resilience. The company should identify, test, and control the security and abuse risks specific to material AI systems, including adversarial manipulation, misuse, data leakage, unsafe tool access, and malicious or unintended system behaviour.

D.2.6 AI Governance in Transactions and Material Change Events. The board should ensure that material transactions and change events include explicit review of AI-related assets, liabilities, dependencies, and governance implications.

Principle D.3 – Data Governance

The board should ensure that data governance is treated as a business responsibility and is sufficient to support the safe, effective, and lawful use of AI.

Provisions

D.3.1 Data Ownership and Quality. Senior accountability for data used in material AI systems should be explicit, and the board should be satisfied that data quality is fit for purpose.

D.3.2 Lawful and Ethical Data Use. Data used in AI should be collected, processed, and retained lawfully and in ways consistent with reasonable stakeholder expectations.

D.3.3 Proprietary Data Assets. Boards should treat proprietary data as a strategic asset and set policies on how company data may be used by third-party models and whether rights to derived improvements are retained.

D.3.4 IP, Provenance and Recordkeeping. The board should ensure that the company can account for the provenance, rights position, and recordkeeping implications of data, models, prompts, outputs, and other artefacts used in or generated through material AI activity.

Section E – People, Culture and Remuneration

AI fundamentally alters the nature of work. A company that deploys AI without deliberate attention to its workforce (its skills, incentives, and trust) will fail to capture the technology's potential and will damage the organisational culture on which long-term success depends.

Principle E.1 – Workforce Impact and Transition

The board should oversee the impact of AI on the workforce and ensure that appropriate provision is made for skills development, redeployment, and workforce transition.

Provisions

E.1.1 Impact Assessment. Management should assess how AI changes roles, skills needs, and workforce composition and should report this as part of the company's people strategy.

E.1.2 Reskilling and Transition. Where AI materially changes role demand or skill requirements, the company should have credible, resourced plans for reskilling, redeployment, or workforce transition.

E.1.3 Workforce Disclosure. Annual reporting should explain how AI affects the workforce, including skills development, support for affected employees, and the company's approach to transition.

Principle E.2 – Remuneration and Incentives

Remuneration structures should support responsible AI adoption and should not create incentives for ungoverned, reckless, or unethical AI use.

Provisions

E.2.1 Outcome-Based Metrics. Where AI is material, incentive structures should measure outcomes enabled by AI rather than outdated activity measures that AI may make irrelevant.

E.2.2 Responsible Use Incentives. Incentive structures should reinforce responsible AI use and allow consequences where poor AI governance destroys value.

Principle E.3 – Culture and the Workforce Contract

The board should ensure that AI is deployed in a manner that preserves and strengthens the trust of the workforce and is consistent with the company's purpose and values.

Provisions

E.3.1 Workforce Engagement. Employees should be told clearly and honestly how AI is being used, how roles may change, and what the company's intentions are before material deployments occur.

E.3.2 Ethical Deployment. AI should not be used in employee surveillance, evaluation, or consequential decision-making in ways that are unexpected, disproportionate, or inconsistent with company values.

E.3.3 Shadow AI. The company should set clear, usable policies for employee use of AI tools, including consumer-grade and personal tools, and support responsible usage with guidance and training.

Section F – Transparency and Accountability

Transparency in AI governance builds the trust of investors, customers, regulators, and employees. Accountability ensures that when AI causes harm or fails to deliver value, there are credible mechanisms for redress. Both are essential to the social licence on which sustainable business depends.

Principle F.1 – AI Governance Disclosure

The board should ensure that the company reports clearly on its approach to AI governance, its material AI activities, and the manner in which AI risk is managed and overseen.

Provisions

F.1.1 Annual Report Disclosure. The annual report should describe board oversight of AI, executive accountability, AI risk management, and material AI risks or incidents.

F.1.2 Material AI Use. Where AI is material to stakeholders, the company should disclose the use in accessible terms and explain the governance arrangements around it.

F.1.3 Consistency with Regulatory Obligations. AI disclosures should be consistent with applicable regulator guidance and with international disclosure obligations where the company operates across borders.

F.1.4 Reliance on AI in Financial Reporting and Board Disclosures. Where AI materially supports financial reporting, management certification, forecasts, narrative reporting, or board-level disclosures, the board should ensure that reliance is controlled, reviewable, and consistent with the company's assurance obligations.

Principle F.2 – Accountability and Redress

The company should maintain credible mechanisms by which stakeholders can raise concerns about AI-driven decisions and seek appropriate redress.

Provisions

F.2.1 Redress Processes. Stakeholders affected by AI-influenced decisions should have accessible routes to raise concerns and seek timely, substantive review.

F.2.2 Explainability. Where AI materially influences consequential decisions affecting individuals, the company should be able to offer a meaningful explanation of the basis for that decision.

F.2.3 Regulatory Cooperation. The company should engage constructively and transparently with relevant regulators regarding AI governance and respond fully to inquiries.

Closing Statement

Good AI governance is not a compliance exercise. It is a discipline that must evolve as rapidly as the technology it governs. The companies that govern AI well will not merely satisfy regulators, they will earn the trust of the customers, employees, and investors on which long-term success depends.

This Code sets a minimum standard of board stewardship for AI. It is not a ceiling. Boards are encouraged to exceed its provisions where their circumstances, values, and ambitions demand it.

The ultimate purpose of AI governance is simply to ensure that organisations retain the capacity to act deliberately, control outcomes, and preserve their agency in a world increasingly shaped by AI.

Further Details

Further information including detailed questions and maturity assessment for Board Directors and Non-Executive-Directors can be found in the book **The Board's Guide to Corporate AI Governance - Stewardship, Strategy, and Accountability in the Age of AI** written by Vincent A. Powell.

ISBN 978-1-9195185-5-8

<https://www.theboardaiguide.com/>