AI in the Boardroom

Strategy, governance, and transformation in the age of AI

Principles for Responsible AI Use by Directors

A standard for AI-enabled board effectiveness that preserves judgement, accountability, and trust.

EIGHT PRINCIPLES · FOUR GOVERNING USE, FOUR MAKING IT WORKABLE

Karim Harbott, CDir, FIoD, FBCS

Chartered Director · Board Adviser · Executive Educator

2026

SUMMARY

The eight principles at a glance

The principles run as a set. The first four govern a director's use of AI itself; the second four are the conditions that make the first four workable. Adopting the first four without the second four produces good intentions with no mechanism; adopting them together produces a standard.

PRINCIPLES 1–4 GOVERNING THE USE

1 **AI augments a director's judgement; it does not replace it**
Form your own view first, then use AI to test it — never to supply it.

2 **Accountability stays with the director**
An AI-assisted decision is simply a decision; the assistance is invisible to the law and the regulator.

3 **Confidential, privileged, and price-sensitive material must be protected**
Board material enters approved, enterprise-grade tools or it enters nothing.

4 **AI outputs must meet professional scepticism**
Fluency is not accuracy; verify before anything is relied upon.

PRINCIPLES 5–8 MAKING IT WORKABLE

5 **The board adopts a written policy for its own use of AI**
Informal practice is not a policy; name an owner and a review cadence.

6 **Directors maintain the AI literacy the role now requires**
Enough understanding to challenge an output rather than defer to it.

7 **AI is used ethically and in line with the organisation's values**
AI-assisted work meets the same standard as any other board conduct.

8 **AI is used where it improves board effectiveness, not for its own sake**
Preparation can be accelerated; deliberation cannot be automated.

INTRODUCTION

Why boards need this standard now

Directors are already using AI. They are using it to prepare for meetings, condense board packs, draft papers, and test their thinking, and much of this use is happening with no policy, no approved tooling, and no shared understanding of what is acceptable. Most boards have built careful oversight of the enterprise's use of AI — a committee that owns it, an inventory of deployments, a stated risk appetite — while leaving the board's own use entirely ungoverned. The result is a quiet accumulation of risk at the most sensitive desk in the organisation: confidential material entering public tools, unverified outputs informing real decisions, and directors carrying exposures they have not examined.

Used well, AI is a genuine aid to board effectiveness. It can help directors absorb more information in less time, prepare more thoroughly, broaden the range of scenarios and counterarguments considered, and free up attention for the work only directors can do: judgement, challenge, and stewardship. The purpose of this standard is therefore not to discourage use. It is to make use deliberate — to replace informal, inconsistent, and sometimes risky habits with a written standard the whole board applies.

The eight principles run as a set. The first four govern the use itself: AI should augment a director's judgement rather than replace it; accountability remains with the director; confidential material must be protected; and outputs must be subject to professional scepticism before they are relied upon. The remaining four are the conditions that make the first four workable: a board-level policy, sufficient AI literacy to challenge rather than defer, use that is ethical and consistent with the organisation's values, and use that serves board effectiveness rather than novelty. Adopting the first four without the second four produces good intentions with no mechanism; adopting them together produces a standard.

Where AI helps a board

AI tools can legitimately support directors, company secretaries, and governance teams across a band of work that is already familiar:

- summarising long board papers, reports, and regulatory developments;
- preparing briefing notes and identifying the questions worth putting to management;
- scanning for industry, competitive, and regulatory signals;
- comparing strategic scenarios and stress-testing options;
- supporting committee preparation and analysis;
- first-pass drafting of minutes, papers, and correspondence; and
- reducing administrative load within the governance function.

None of this is exotic, and that is the point: the technology is already doing ordinary board work. What AI lacks is experience, context, accountability, and fiduciary responsibility. Every use case above is an act of preparation or analysis; none is a decision. The principles that follow exist to keep that line intact.

GOVERNING THE USE

1 AI augments a director's judgement; it does not replace it

WHY THIS MATTERS

The value of a non-executive director is judgement formed by experience, applied independently. Delegating that thinking to a model is an abdication of the responsibility a director is appointed to exercise, and habitual reliance carries a slower harm: the director's own faculty atrophies precisely as it is needed most. The distinction that keeps augmentation on the right side of this line is one of order. A director who forms a view first and then uses AI to test it — to find the gap, the blind spot, the counter-argument not yet considered — is using the tool to check judgement. A director who lets the machine think first is using it to supply judgement, and that is where displacement begins.

PRACTICAL APPLICATION

Use AI freely for preparation, research, and synthesis. For matters of judgement, form your own position before consulting the tool, then use it adversarially.

DO

+ Form your own view first, then use AI to identify its weaknesses.
+ Ask AI for the counter-argument, the missing scenario, and the question you have not asked.
+ Use AI to prepare more thoroughly, so board time goes to challenge and deliberation.
+ Bring your own conclusions to the boardroom, not the tool's.

DON'T

– Ask AI what to decide, then look for reasons to agree.
– Substitute an AI summary for reading the paper on any matter you will vote on.
– Let AI-drafted questions replace the ones your own experience would raise.
– Allow the tool's convenience to shorten your own thinking.

GOVERNING THE USE

2 Accountability stays with the director

WHY THIS MATTERS

A director's accountability is fixed by law and does not move because a model helped prepare the analysis. Under the Companies Act, a director must exercise reasonable care, skill, and diligence; in regulated financial firms, individual accountability is allocated under the Senior Managers and Certification Regime. An AI-assisted decision is, for these purposes, simply a decision — the assistance is invisible to the standard and to the regulator. A decision resting on an unverified AI summary that proves wrong on a material fact is the director's to answer for, not the tool's.

PRACTICAL APPLICATION

Treat AI output as you would any external input relied upon in a decision: verified in proportion to the significance of the matter, and never assumed to have been checked by someone else.

DO

+ Verify material facts, figures, and references before relying on them.
+ Scale verification to the stakes: a briefing summary and an M&A analysis do not deserve the same scrutiny.
+ Use more than one source where the decision is significant.
+ Ensure the board minutes would survive the disclosure that AI contributed to the analysis.

DON'T

– Assume an AI output has been validated by anyone.
– Rely on AI-generated legal, financial, or regulatory interpretation without professional review.
– Treat the tool's contribution as diluting your responsibility for the outcome.
– Let convenience substitute for the diligence the role requires.

GOVERNING THE USE

3 Confidential, privileged, and price-sensitive material must be protected

WHY THIS MATTERS

Boards handle the organisation's most sensitive information: strategy, transactions, litigation, people, and, in listed companies, inside information. A board pack pasted into a public AI tool may be retained, used for training, or exposed to third parties, and it is a disclosure the director may struggle to explain afterwards. This is also where the central trade-off of board AI use sits: the organisational context that makes AI genuinely useful is the same confidential material that creates the exposure. Usefulness and exposure rise together, which is why the answer is approved, enterprise-grade tooling rather than abstinence.

PRACTICAL APPLICATION

Know how any tool you use processes, retains, and trains on data before board material goes anywhere near it. Where an approved enterprise tool exists, use it; where none exists, that is a gap for the board to close, not a licence to improvise.

DO

+ Use the organisation's approved, enterprise-grade tools for any board-related work.
+ Understand the retention and training terms of every tool you use.
+ Treat board papers, legal advice, and inside information as barred from unapproved tools without exception.
+ Raise the absence of an approved tool with the chair or company secretary rather than working around it.

DON'T

– Paste board packs, draft results, or transaction material into public AI tools.
– Assume a paid consumer subscription offers enterprise-grade protection.
– Rely on deleting a conversation to undo a disclosure.
– Ignore data protection, market abuse, and confidentiality obligations because the interface feels informal.

GOVERNING THE USE

4 AI outputs must meet professional scepticism

WHY THIS MATTERS

AI systems produce output that can be inaccurate, incomplete, biased, outdated, or invented, and they present it all with the same fluent confidence. Fluency is not accuracy. A hallucinated figure or a quietly biased analysis that enters the boardroom unchallenged becomes part of a decision the board owns in full. Directors already know how to handle confident, unaccountable input — they challenge management papers and consultant reports for a living — and AI output deserves the same treatment, applied more strictly because no author is answerable for it.

PRACTICAL APPLICATION

Interrogate AI output as you would a paper from an adviser with no liability for being wrong: check the load-bearing facts, test the logic, and ask what is missing.

DO

+ Cross-check any figure, citation, or claim on which a decision will rest.
+ Ask what context the tool cannot have had, and what that omission changes.
+ Probe whether the output reflects your organisation's actual position or a generic one.
+ Use AI to open lines of inquiry, then pursue them with primary sources.

DON'T

– Confuse a confident, well-written answer with a correct one.
– Accept citations or quotations without verifying they exist and say what is claimed.
– Assume the output is neutral; models carry the biases of their training data.
– Treat plausibility as evidence.

MAKING IT WORKABLE

5 The board adopts a written policy for its own use of AI

WHY THIS MATTERS

Where no policy exists, every director sets their own standard, and the true position of the board is the practice of its least careful member. Ungoverned use collects exactly where the stakes are highest — real board business, public tools, no verification discipline — and the exposure is live now, whatever the board eventually decides about where to draw its lines. A written policy is the cheapest form of governance available: it converts eight good intentions into a standard that can serve as induction material, a reference in a dispute, and a shared expectation around the table.

PRACTICAL APPLICATION

Adopt a short policy covering directors, committees, and the governance function. Name its owner — usually the board together with the company secretary — and set a review cadence that matches the pace of the technology.

DO

+ Define acceptable and prohibited uses in terms a director can apply unaided.
+ Name the approved tools and the material that must never leave them.
+ State the verification expected before AI output informs a board decision.
+ Review the policy at least annually, and sooner when tools or regulations move.

DON'T

– Assume informal practice is a policy.
– Tolerate different standards across directors or committees.
– Treat directors' own use of AI as outside the scope of governance oversight.
– Wait for an incident to supply the motivation.

MAKING IT WORKABLE

6 Directors maintain the AI literacy the role now requires

WHY THIS MATTERS

A director cannot apply professional scepticism to a technology they do not understand at a working level. Literacy here does not mean technical depth; it means knowing what the tools can and cannot do, how they fail, and what their use exposes — enough to challenge an output rather than defer to it, and enough to govern the enterprise's AI with credibility. A board that lacks this literacy does not avoid AI risk; it simply stops seeing it.

PRACTICAL APPLICATION

Treat AI literacy as part of a director's continuing professional development, supported by the board rather than left to individual initiative.

DO

+ Build a working understanding of what current tools do well and where they fail.
+ Use the tools yourself, within policy, so the understanding is practical rather than reported.
+ Participate in structured director education on AI and its governance.
+ Keep abreast of developments material to the board's oversight duties.

DON'T

– Treat AI literacy as a management competence the board can borrow when needed.
– Rely entirely on others' interpretations of what the technology means for the organisation.
– Overestimate reliability because the output reads well.
– Disengage because the subject appears technical; the duty of skill does not carve it out.

MAKING IT WORKABLE

7 AI is used ethically and in line with the organisation's values

WHY THIS MATTERS

AI output can carry bias into board discussion invisibly: in whose perspectives a summary amplifies, in what a scenario analysis treats as normal, in which risks a briefing quietly ranks below others. A board that would never accept a conflicted adviser should not accept an unexamined one. Directors' own use of AI also sets the tone for the organisation: a board that cuts corners with the technology forfeits its standing to demand that management not do so.

PRACTICAL APPLICATION

Apply the same ethical standard to AI-assisted work as to any other board conduct, and stay alert to the ways the tool's framing can shape the board's.

DO

+ Consider whose interests and perspectives an output may under-represent.
+ Hold AI-assisted work to the organisation's values and governance standards.
+ Notice when the tool's framing is steering the discussion, and name it.
+ Keep use proportionate to purpose.

DON'T

– Assume AI output is neutral because it is machine-generated.
– Dismiss ethical questions because the task looks administrative.
– Use AI in ways that would embarrass the board if disclosed.
– Model behaviour that the board would sanction elsewhere in the organisation.

8

MAKING IT WORKABLE

AI is used where it improves board effectiveness, not for its own sake

WHY THIS MATTERS

The test for any AI use in board work is whether it improves the quality, speed, or breadth of governance — better-prepared directors, sharper challenge, wider scenarios, faster cycle times in the governance function. Use that fails this test adds cost and risk without return: new confidentiality exposure, new verification burden, new complexity in processes that were working. Deliberation is the clearest boundary. Preparation can be accelerated; the discussion through which a board reaches its decision cannot be automated without ceasing to be governance.

PRACTICAL APPLICATION

Adopt uses that demonstrably improve board effectiveness, experiment deliberately within the policy's guardrails, and drop uses that do not earn their place.

DO

+ Prioritise uses with a clear return on effectiveness: preparation depth, breadth of analysis, and speed of the governance function.
+ Pilot new uses deliberately and review whether they delivered.
+ Encourage experimentation within the policy's boundaries.
+ Retire uses that add complexity without benefit.

DON'T

– Adopt tools because peers or fashion suggest it.
– Automate any part of the board's actual deliberation or decision-making.
– Add AI steps to governance processes that were not the constraint.
– Treat adoption itself as the objective.

CONCLUSION

The ordinary governance reflex, applied to the board's own desk

AI will become a standard part of how directors prepare, analyse, and challenge, in the way that search and spreadsheets did before it. What it will not become is a substitute for the things a board exists to provide: judgement, accountability, stewardship, and the willingness to challenge. The eight principles in this standard are the ordinary governance reflex — decide, document, assign ownership, review — applied to the one surface most boards have so far exempted from it: their own desk. A board can adopt them now, name an owner, and revisit them as the technology moves. The alternative is the current default, in which the use happens anyway, ungoverned, and the exposure sits with each director alone.

ABOUT THE AUTHOR

Karim Harbott

CDir, FIoD, FBCS · Chartered Director, Board Adviser & Executive Educator

Karim Harbott is a Chartered Director, board adviser, and executive educator who helps boards turn AI from a source of anxiety into a source of advantage. He serves as a senior independent and non-executive director, and teaches strategy and AI governance to directors through the Institute of Directors and the Saudi Centre for Governance. After two decades leading enterprise transformation — at McKinsey, Barclays, and the consultancy he co-founded — he now works where strategy, governance, and AI meet, and where most boards are least confident. He is the author of *The 6 Enablers of Business Agility* and writes *AI in the Boardroom*.

AI in the Boardroom

Strategy, governance, and transformation in the age of AI

READ & SUBSCRIBE
aiintheboardroom.com

LINKEDIN
linkedin.com/in/karimharbott

WEBSITE
karimharbott.com

© Karim Harbott 2026 · This document may be shared in full and unaltered for board and governance use.