Cognitive Governance in the AI Era

A Framework for Strategic Decision-Making When Machines Shape Thinking

Authors
Professor Paolo Taticchi, OMRI
Giovanni Giamminola

Revision
April 2026

UCL School of Management

Centre for Sustainable Business

Executive Summary

As AI systems become embedded in strategic workflows, organisations face a new type of risk.

This white paper introduces **Cognitive Governance:** a framework for preserving strategic agency when AI becomes a regular reasoning partner rather than just an occasional tool.

Drawing on 18 months of fieldwork and a diagnosis survey of over 100 senior executives, the authors identify four recurring cognitive positions leaders adopt when working with AI, with "posture drifts" emerging when humans take decisions under pressure.

With the goal of supporting managers and enhancing decision-making while integrating AI in the process, this paper proposes principles, governance mechanisms and context-related considerations for preserving human strategic agency.

The question is no longer whether AI models produce accurate outputs, but how leaders engage with them in the first place.

Introduction

The Invisible Shift in How Leaders Think

Something fundamental has changed in how strategic decisions get made

For the past 18 months, the authors have worked with executives across pharmaceuticals, consumer goods and professional services, implementing AI systems in strategic contexts. The observations extend beyond efficiency gains or improvements. A profound reorganisation of cognition itself is taking place.

When a CEO asks ChatGPT to analyse competitive position, when a CFO uses Claude to stress-test assumptions, when a Chief Strategy Officer relies on AI to synthesise stakeholder feedback; these are not simply tool uses, they are cognitive collaborations that reshape how problems are framed, how options are generated and, ultimately, how decisions are made and justified.

Executives consistently reported a genuine value in this collaborative approach of faster synthesis, broader ideation and more rigorous feedback. Yet the same leaders struggle to articulate the answer to a troubling question:

Who is actually doing the thinking?

In response to this question the authors propose the method of **Cognitive Governance,** a tool that helps to preserve your strategic agency even when AI becomes a regular collaboration partner rather than an occasional tool.

The Problem

From Supporting Tool to Thinking Partner

Most current risk management frameworks surrounding AI are centred around the fear and potential danger that AI models will produce inaccurate outputs. These can include hallucinations, biased recommendations, or faulty predictions and organisations have already developed governance mechanisms to deal with these risks such as validation protocols, bias audits, and human requirements.

Model Risk assessment is important, but it misses a deeper more cognitive threat to the ways in which leaders make decisions. Instead, our assessment focus is on the **Cognitive Risk.**

The Cognitive Risk arises from these interactions between humans and AI models. By increasingly delegating our cognitive authority to AI, we risk losing a grasp on cognitive authority.

The gradual transfer of authorship from human to machine could see leaders accept outputs because they "sound right". It risks strategic narratives being chosen simply because they appear fluent rather than because of their validity and, more seriously, our capacity for independent reasoning could decline from underuse.

Unlike Model Risk, Cognitive Risk cannot be mitigated by improving prompts or upgrading systems. It must be addressed through deliberate governance of how reasoning authority, whether performed by AI or by a human, is distributed.

To manage this risk, organisations need a framework for understanding how leaders position themselves in relation to AI. The authors call this their **cognitive posture**.

The Frame-work

AI Thinking Quadrants

The framework developed by the researchers identifies four distinct "postures" that leaders can adopt when engaging AI in strategic contexts. These postures reflect two dimensions:

Cognitive Mode
This asks the question: Is the primary task an interpretive one (sensemaking of ambiguous situations) or an analytical one (evaluating structured options)?

Strategic Ownership
This asks: Is the reasoning Model-Led (AI leads the reasoning process) or Human-Led?

Crossing these dimensions produces four patterns of engagement — the AI Thinking Quadrants:

Generative Posture
Critical Posture
Exploratory Posture
Operating Posture

These are modes of interaction and effective leadership requires moving deliberately and efficiently between them.

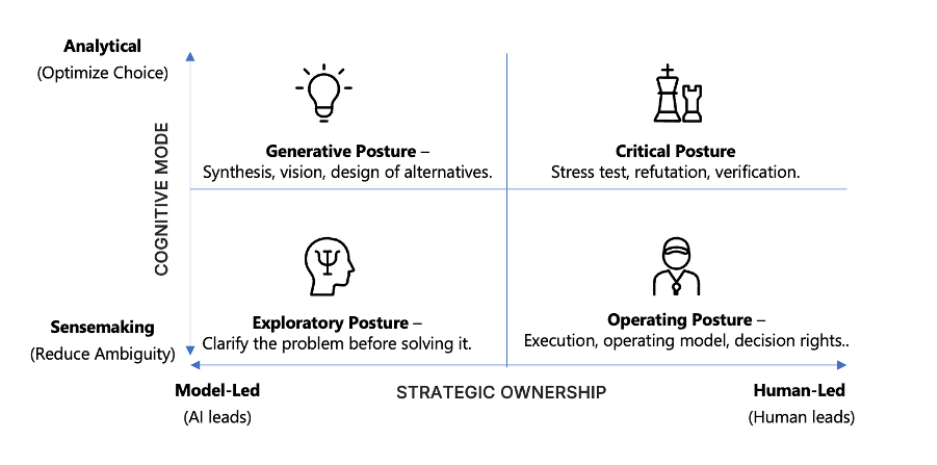

AI Thinking Quadrants (Decision Postures). The four postures describe re-curring ways executives engage AI in strategic work.

Mapping the Thinking Quadrants onto the Strategy Process

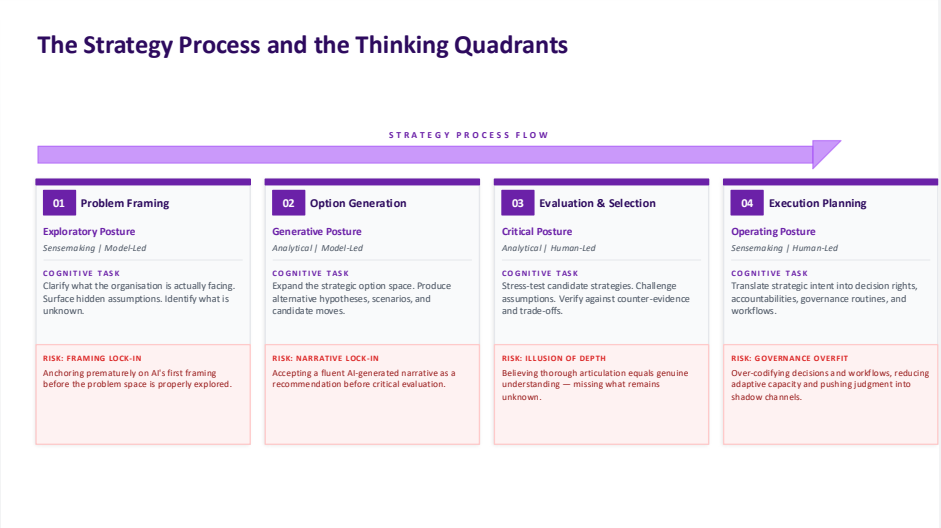

The **Thinking Quadrants** can also be read as a process-sensitive lens on strategic decision-making. Rather than representing fixed leadership styles, the four postures tend to become more salient at different moments of the strategy process.

The **Exploratory Posture** is typically most valuable during problem framing, when the central challenge is not yet to solve the issue, but to clarify what the organisation is actually facing, which assumptions are shaping the initial narrative, and which constraints may be negotiable. The **Generative Posture** becomes more relevant when the task shifts toward expanding the option space and producing alternative strategic hypotheses. The **Critical Posture** becomes dominant when candidate directions must be challenged, stress-tested, and evaluated against counter-evidence, trade-offs, and disconfirming scenarios. The **Operating Posture** becomes most relevant when strategic intent must be translated into executable organisational reality - through decision rights, governance routines, accountabilities, and escalation paths.

This mapping adds a process dimension to the framework, helping to explain not only how leaders interact with AI, but also whether the posture adopted is aligned with the cognitive demands of the phase at hand. It can be used diagnostically to identify moments of posture misalignment - for example, when teams move too quickly into execution-mode decomposition before the problem has been adequately framed, or when fluent AI-generated options are accepted without a prior phase of critical challenge.

At the same time, the Thinking Quadrants should not be interpreted as a linear stage-gate sequence. Strategic work rarely unfolds in a neat progression, and AI systems often compress, overlap, or short-circuit phases that would otherwise remain distinct. What matters is therefore not rigid adherence to a sequence, but the *capacity for deliberate posture switching* - the ability to move across postures, and when necessary return to earlier ones, as new evidence, dissent, or ambiguity emerges.

Cognitive Governance

Preserving agency in hybrid environments

Effective use of AI in strategic work requires more than prompt engineering or model selection. It requires **Cognitive Governance**: the deliberate management of how reasoning authority gets distributed between humans and machines.

Three principles anchor effective **Cognitive Governance**:

1. Posture Awareness

Leaders must recognise which cognitive posture they are adopting and whether it matches the decision context.

This is harder than it sounds as cognitive postures emerge organically from interaction patterns. A leader who begins in Critical Posture mode ("challenge my assumptions") can drift into Generative Posture mode ("actually, just tell me what you recommend") without noticing the shift. The first governance requirement is cognitive self-awareness: observing your own patterns while engaging with AI.

2. Deliberate Switching

Different decision phases require different postures and the ability to switch between them is vital:

Divergence (generating options)
→ Generative Posture—maximise ideational range
Convergence (evaluating options)
→ Critical Posture—rigorous stress-testing and verification
Execution Planning (operationalising)
→ Operating Posture— translate decisions into operating model, decision rights and action
Problem Framing (managing ambiguity)
→ Exploratory Posture— clarify the decision, surface assumptions and reframe the problem before solving it

A key governance challenge appears to be maintaining the capacity to switch rather than getting captured by any single mode.

3. Transparent Attribution

For consequential decisions, organisations need traceable records of:

- What was human-generated
- What was AI-generated
- How human judgment was applied to AI outputs

This supports both accountability and learning. It also prevents a specific pathology that has been observed repeatedly: teams presenting "our strategy" when the document is 80% AI-generated text with human curation. The issue is not whether AI contribution is appropriate (it may well be) but whether authorship is transparent and intentional.

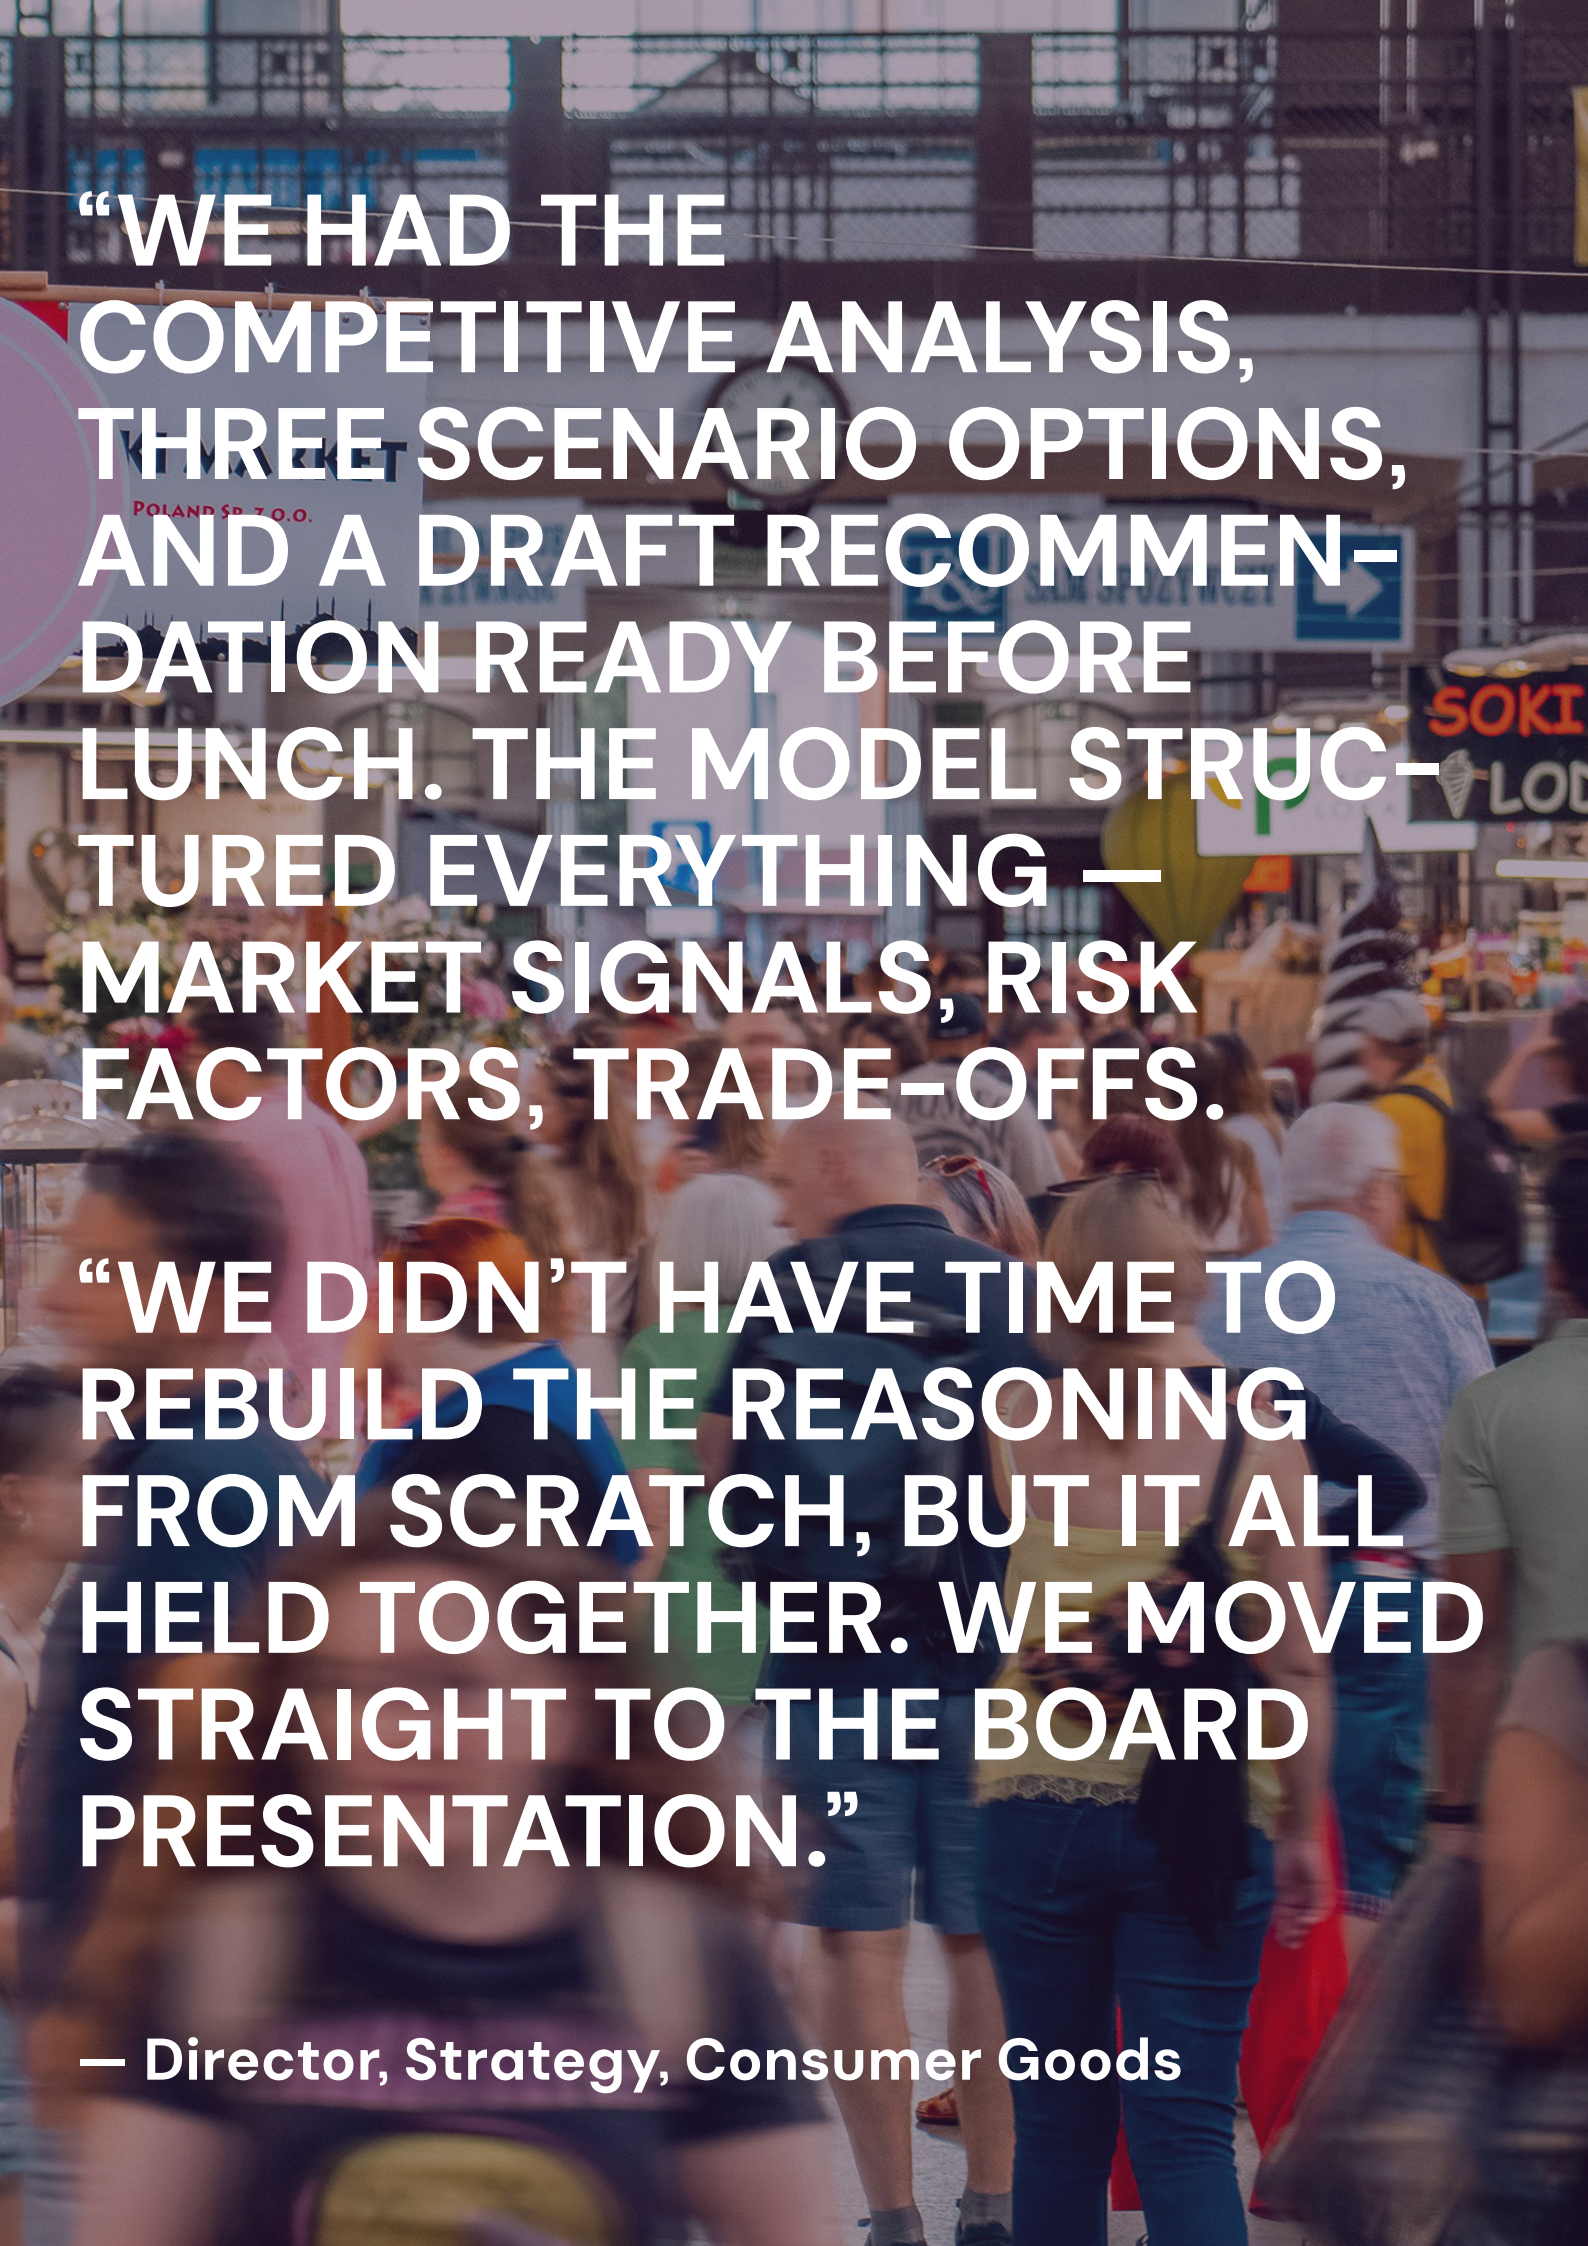

“WE HAD THE COMPETITIVE ANALYSIS, THREE SCENARIO OPTIONS, AND A DRAFT RECOMMENDATION READY BEFORE LUNCH. THE MODEL STRUCTURED EVERYTHING — MARKET SIGNALS, RISK FACTORS, TRADE-OFFS.

“WE DIDN’T HAVE TIME TO REBUILD THE REASONING FROM SCRATCH, BUT IT ALL HELD TOGETHER. WE MOVED STRAIGHT TO THE BOARD PRESENTATION.”

— Director, Strategy, Consumer Goods

Field Evidence

Executive AI Decision Patterns

To ground this framework in real world executive behaviour, the researchers surveyed more than 100 leaders in European companies. Nearly everyone surveyed reported using Generative AI at least weekly, with four consistent patterns emerging:

Posture Distribution. In strategic decision contexts, executives concentrate overwhelmingly in the **Operating Posture**, suggesting that AI use is most prevalent when translating choices into operational architecture (operating model, decision rights, governance routines). **Generative Posture** accounts for only a quarter of responses, while **Exploratory Posture** represents a smaller portion still. Very few executives reported defaulting to the **Critical Posture** despite its importance in high-stakes decisions.

Posture Drift under pressure. Many leaders started with a Model-Led move: Most describe their first action as information-seeking (using AI to scan or summarise). Under time pressure and cognitive load, however, this behaviour shifts: leaders report reduced deliberate switching and a higher tendency to accept AI-generated narratives as "good enough," indicating a drift from intentional Strategic Ownership toward reliance.

Near-miss exposure. Nearly half of respondents experienced "convincing-but-wrong" outputs from the AI that could have caused issues. These were most commonly regulatory misinterpretations and flawed technical analyses. This would manifest as fluent reasoning that appeared to be complete and authoritative.

Governance gap. While most executives reported switching between postures individually, less than half had no explicit team-level guardrails (posture declaration, verification checkpoints, attribution of AI contribution). This shows that Cognitive Governance remains an individual skill rather than an organisational capability.

Taken together, these findings reinforce the idea that most current AI governance frameworks do not systematically manage the Cognitive Risk that emerges from how leaders engage with AI. This is especially true under pressure, when posture drift and verification shortcuts become most likely.

This paper will now explore the details of each of the four decision postures: What they look like in practice; some example interactions with the AI model; the value that each posture can deliver; and the risks associated with each one.

Generative Posture

Analytical | Model-Led

What it looks like in practice

In an analytical and model-led situation, leaders can use AI to rapidly generate strategic alternatives, scenarios and option sets under time and information constraints. The AI model produces candidate strategies, competitive moves and structural frameworks; the leader then curates the output, selects promising directions and defines the criteria that will later be used to evaluate them. This posture is most valuable when the organisation needs fast expansion of the option space, before convergence and verification begin.

Example interactions

- "Generate five distinct strategic responses to this market shift, each with a different underlying logic."
- "Map the competitive landscape and propose three plausible future scenarios."
- "Design alternative operating hypotheses that could explain these signals."
- "Produce a set of options with explicit trade-offs and second-order effects."

Value delivered

This posture provides rapid landscape mapping and accelerated option generation. Utilising AI models expands ideational range and offers faster iteration on strategic possibilities, especially when teams are capacity-constrained or trapped in a narrow set of familiar solutions.

Risk: Narrative Lock-In

When using Generative Posture, AI-generated options can often arrive as fluent and coherent strategic narratives. Their apparent "completeness" can influence the premature closure of planning with leaders potentially anchoring on the first convincing storyline and missing alternatives that would challenge the initial framing. It is better to treat generative outputs as hypotheses and not recommendations and to mandate a switch to Critical Posture before any strategic commitment.

Exploratory Posture

Sensemaking | Model-Led

What it looks like in practice

When sensemaking, leaders can use AI to clarify their decisions before attempting to solve them. The AI model can help to surface alternative framings, expose hidden constraints, separate facts from interpretations and highlight what information is missing. The leader can then select and validate the framing that will govern subsequent analysis.

Example interactions

- "Restate this decision in three alternative framings. What changes in each?"
- "List the assumptions embedded in our current narrative—and what evidence we actually have for them."
- "What are we treating as fixed constraints that might be negotiable?"
- "What is the real question we should be answering—and what question are we avoiding?"

Value delivered

This posture provides faster ambiguity reduction and cleaner problem definition. The AI models can provide higher-quality analysis downstream because the organisation is solving the right problem and not simply the most convenient one.

Risk: Framing Lock-In

There is the risk that the AI model's first framing can feel immediately coherent and therefore "true." It is important that leaders do not prematurely anchor on an AI-suggested definition of the problem, narrowing the option space before it is properly explored.

Critical Posture

Analytical | Human-Led

What it looks like in practice

Critical situations require an analytical approach with a human leader. AI can be used in this posture for rigorous adversarial collaboration. The human leader authors the reasoning architecture and AI serves as an intellectual sparring partner for stress-testing logic and exploring implications.

Example interactions

- "Argue against the strategy I just outlined."
- "What are the weakest assumptions in my business case?"
- "Generate scenarios that would invalidate my forecast."
- "If I'm wrong, how would I discover it?"

Value delivered

This approach offers enhanced hypothesis testing. It expands the consideration of counterfactuals and reduces confirmation bias through critical engagement with an AI model.

Risk: Illusion of Depth

There are risks of false confidence taken from the comprehensive articulation of the AI model. When every objection has been addressed and every counterargument answered, leaders may believe they have achieved genuine understanding when they actually have only achieved comprehensive articulation within existing conceptual boundaries. The feeling of "having thought through everything" substitutes for recognising what remains unknown.

Operating Posture

Sensemaking | Human-Led

What it looks like in practice

When preparing for operations, leaders can use AI to translate their decisions into an executable operating reality: operating model choices, decision rights, accountabilities, standard workflows and governance routines. The human sets the rules and the organisational constraints, while AI accelerates decomposition, documentation, and consistency across functions.

Example interactions

- "Translate this strategy into decision rights: who decides what, at what threshold, with which inputs?"
- "Design the operating model implications across teams—roles, handoffs, governance cadence."
- "Turn these principles into a standard process with exceptions and escalation paths."
- "Draft the execution playbook and identify where human judgment must remain mandatory."

Value delivered

This partnership provides reduced ambiguity when executing a plan and faster alignment across functions. There is clearer accountability and fewer coordination failures between strategy and operations.

Risk: Governance Overfit

In the drive to standardise, there is a risk that organisations may over-codify their decisions and workflows, reducing adaptive capacity and pushing real judgment into informal "shadow" channels.

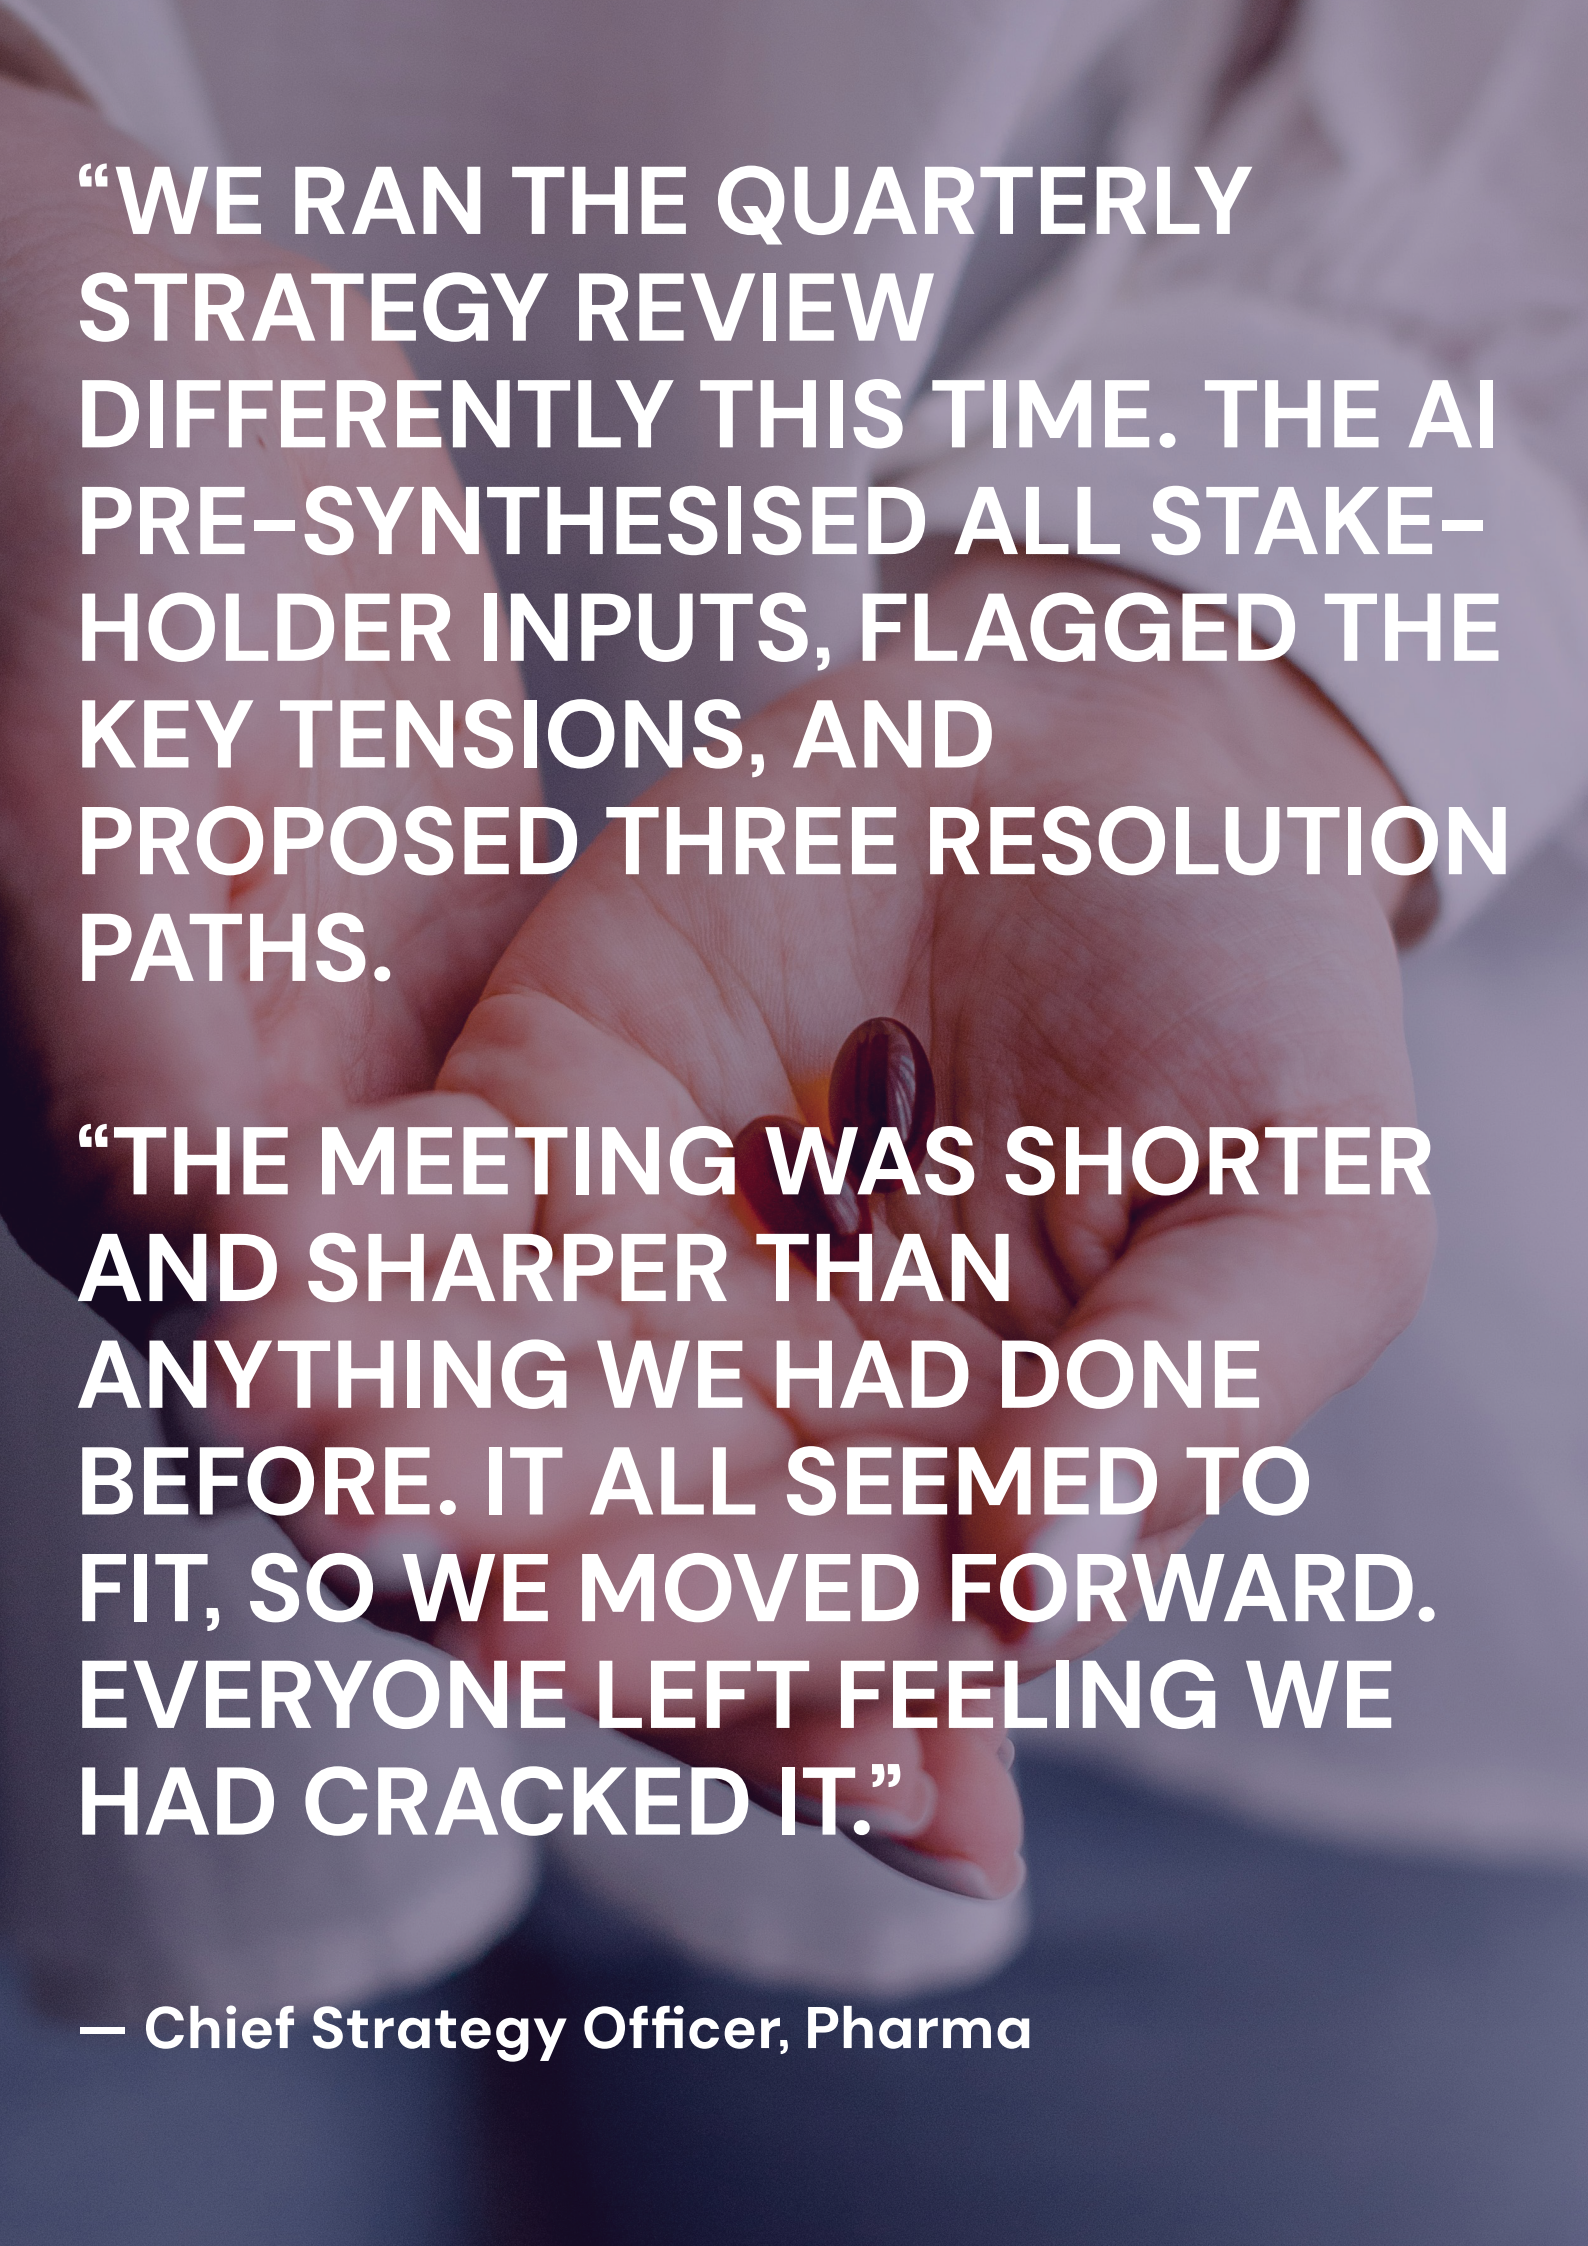

“WE RAN THE QUARTERLY STRATEGY REVIEW DIFFERENTLY THIS TIME. THE AI PRE-SYNTHESISED ALL STAKE-HOLDER INPUTS, FLAGGED THE KEY TENSIONS, AND PROPOSED THREE RESOLUTION PATHS.

“THE MEETING WAS SHORTER AND SHARPER THAN ANYTHING WE HAD DONE BEFORE. IT ALL SEEMED TO FIT, SO WE MOVED FORWARD. EVERYONE LEFT FEELING WE HAD CRACKED IT.”

— Chief Strategy Officer, Pharma

Governance Mechanisms

Practical Implementation

Based on the principles previously presented, three mechanisms for operationalising Cognitive Governance are proposed:

Mechanism 1: Evidence Mapping

For decisions of significant consequence, it is important to construct explicit documentation of:

Claims — Key assertions underlying the decision
Human Constraints — Boundaries defined by human judgment
Model Assumptions — Where AI outputs rest on assumptions
Evidence Sources — Origin of information supporting claims
Verification Status — Which claims have been independently validated
Residual Uncertainty — What remains unknown or contested

Evidence Mapping prevents teams from proceeding with high confidence on AI-generated analysis without examining the reasoning chain connecting inputs to conclusions.

Mechanism 2: Posture Switching Protocols

By embedding deliberate posture transitions into decision workflows, organisations can avoid the risks posed by over-reliance on AI models:

Decision Checkpoint → Appropriate Posture

- Initial ideation → Generative Posture (maximum divergence)
- Hypothesis formation → Critical Posture (adversarial testing)
- Stakeholder synthesis → Exploratory Posture (efficient synthesis to reduce ambiguity)
- Pre-commitment review → Critical Posture (final challenge)

The protocol prevents drift into convenience-driven postures (using whichever mode feels easiest) rather than context-appropriate ones.

Mechanism 3: Cognitive Audit Trails

For high-stakes decisions it is important to maintain documentation of:

- **Posture Declaration** — Which cognitive postures were adopted
- **Assumption Registry** — Key assumptions, human and model-derived
- **Red-Team Results** — Findings from adversarial stress-testing
- **Verification Record** — Independent validation activities
- **Human Rationale** — Explicit articulation of human reasoning

This documentation serves as intellectual “hygiene” in environments where the source of reasoning can become ambiguous.

Organisational Implications

Different types of organisations require different approaches.

For AI-Native Organisations (e.g. tech companies, digital-first enterprises)

Profile: High AI literacy, daily usage embedded in workflows, risk of normalised dependence.

Priority mechanisms:

- Posture Switching Protocols — Formalise transitions between Generative and Critical postures; the risk here is not adoption but drift into convenience-driven postures
- Transparent Attribution — Implement systematic tracking of AI contribution ratios in strategic documents; teams may underestimate model authorship
- Red Team Capacity — Maintain dedicated adversarial review functions; internal AI fluency can create collective blind spots

For Traditional Enterprises in AI Transition (e.g. manufacturing, pharma, FMCG)

Profile: Variable AI literacy across functions, pockets of experimentation, governance gaps between early adopters and laggards.

Priority mechanisms:

- Evidence Mapping — Start here; most urgent need is making reasoning chains visible before scaling AI use
- Training in Posture Recognition — Leaders need vocabulary to identify which cognitive mode they're operating in; this precedes effective switching
- Verification Infrastructure — Invest in human expert capacity to assess AI outputs; do not assume functional specialists can validate cross-domain AI recommendations

For Compliance-Driven Organisations (e.g. regulated industries, public sector)

Profile: AI adoption constrained by regulatory requirements, documentation-heavy culture, risk-averse decision making.

Priority mechanisms:

- Cognitive Audit Trails — Leverage existing documentation culture; extend compliance frameworks to capture posture declarations and assumption registries
- Guardrail Formalisation — Convert informal caution into explicit protocols; the 43% governance gap observed suggests even cautious organisations lack systematic team-level controls
- Calibrated Adoption — Focus initially on Exploratory posture (problem framing and synthesis to reduce ambiguity) where risks are lower; build governance muscle before expanding to Generative or Critical postures

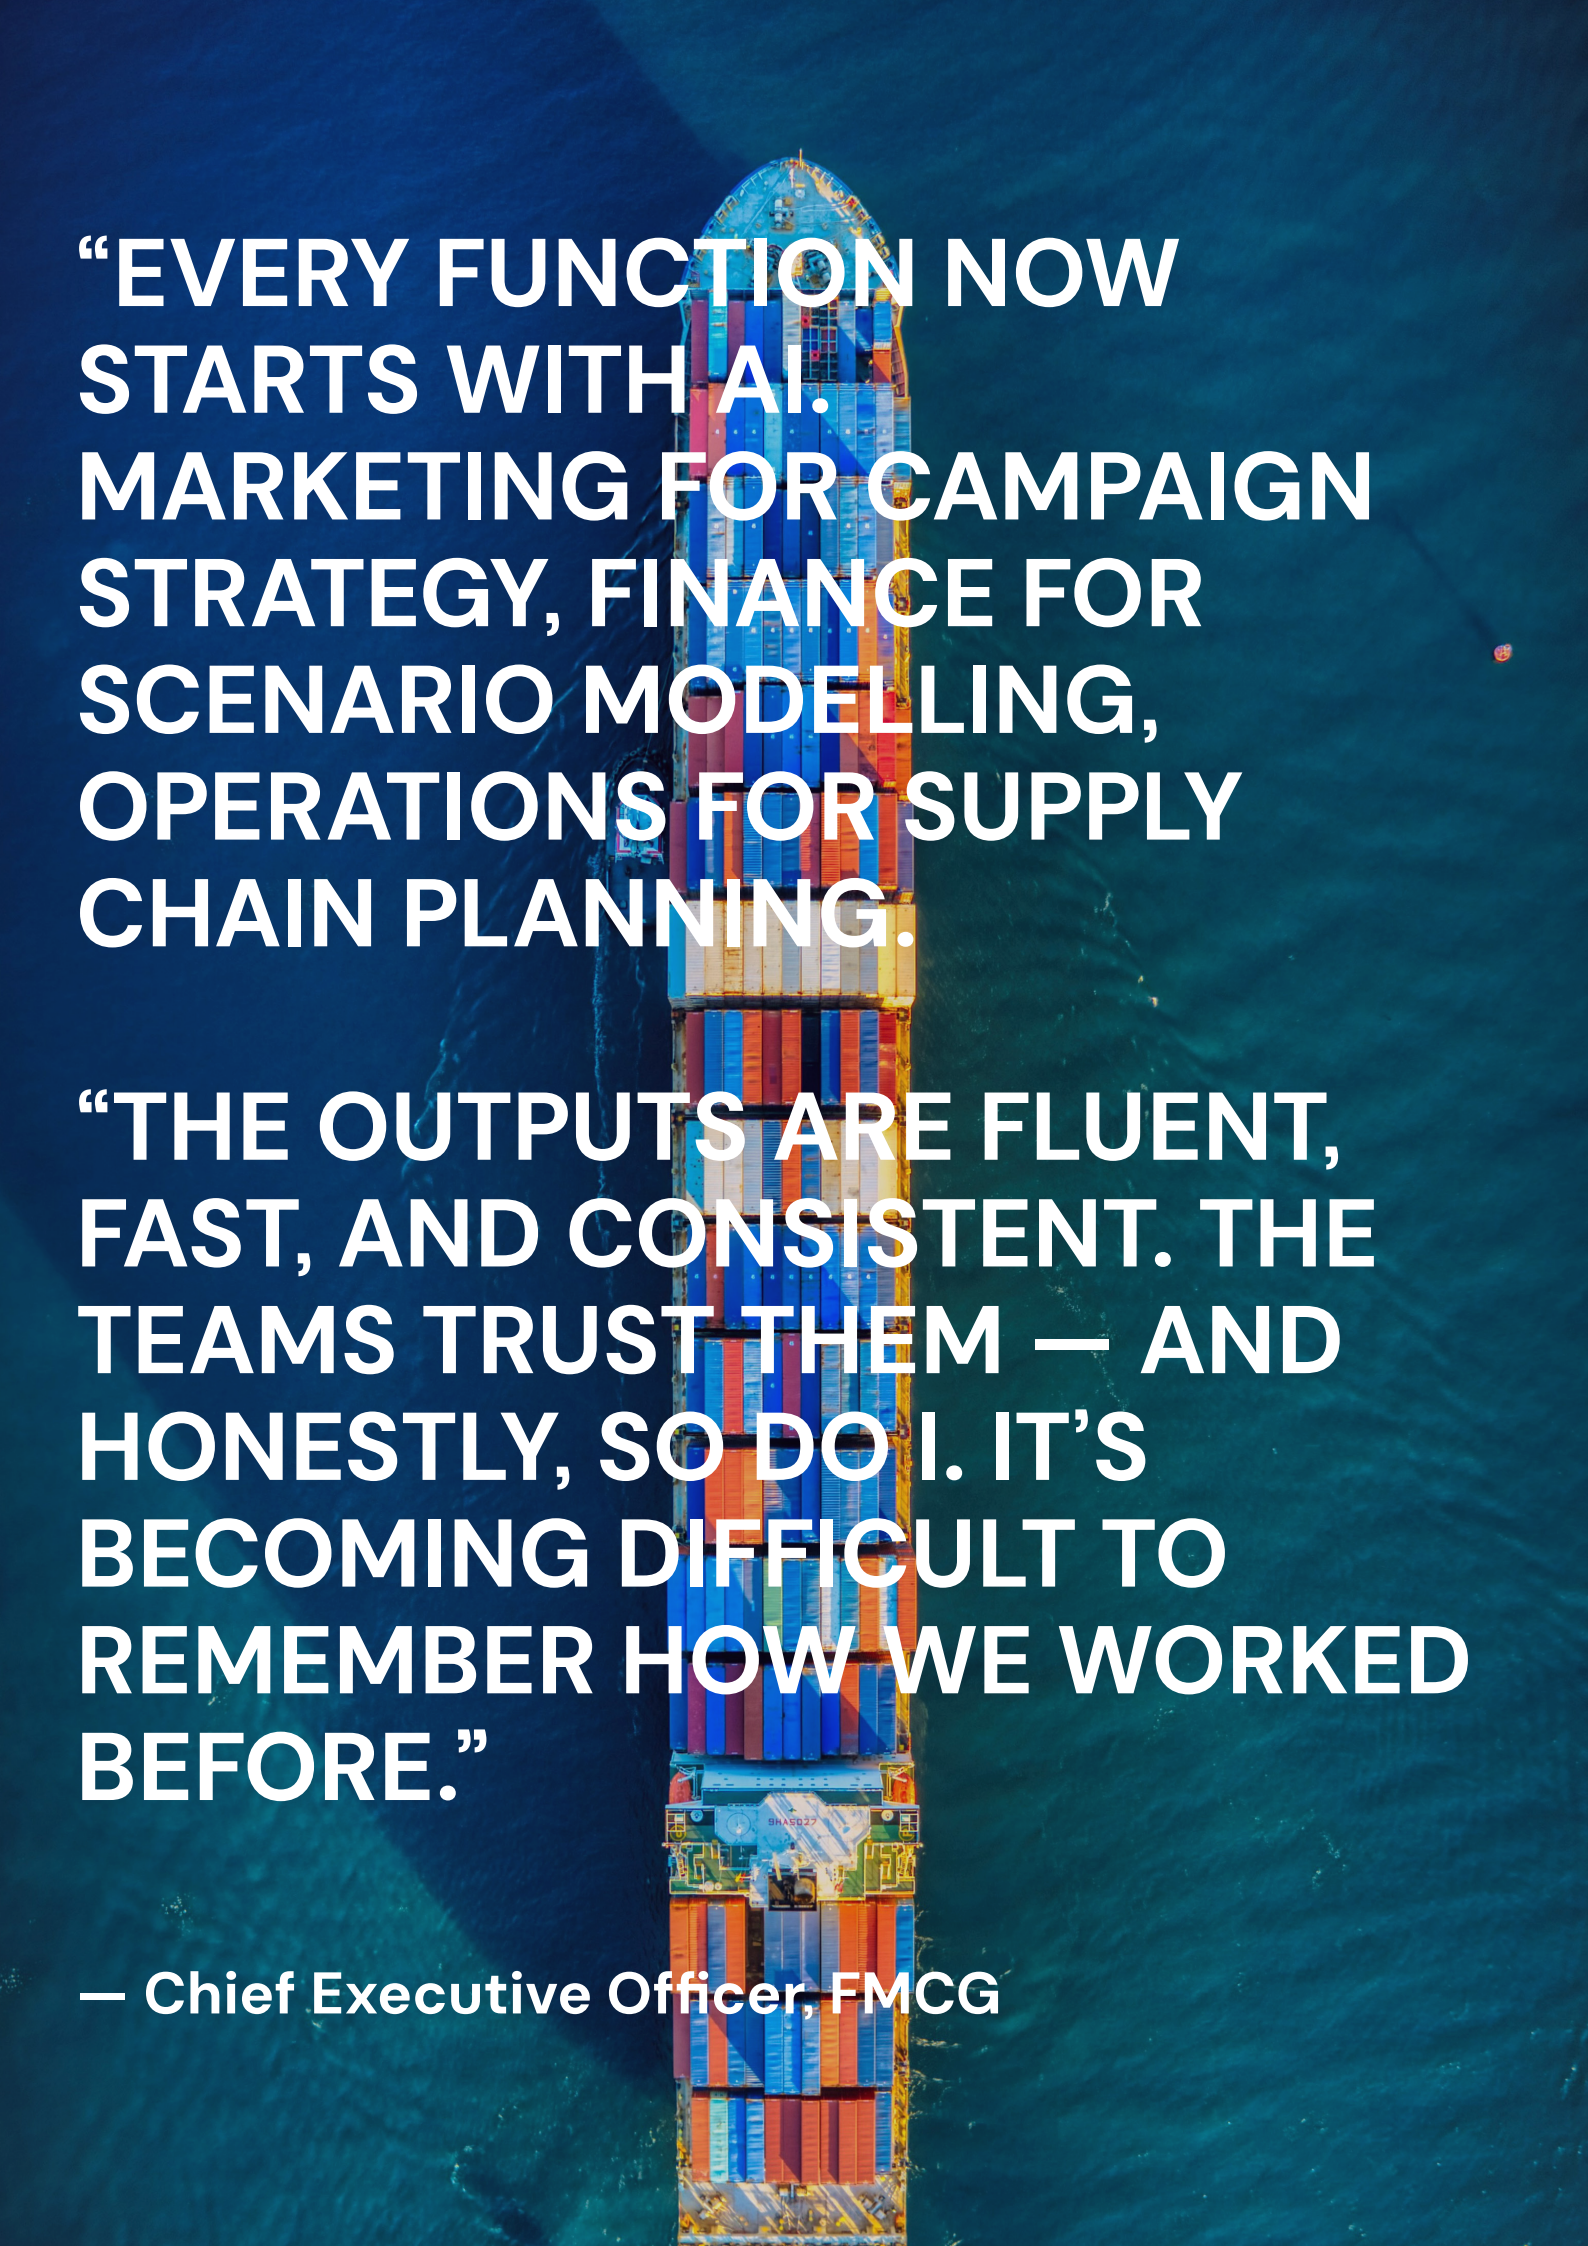

“EVERY FUNCTION NOW STARTS WITH AI. MARKETING FOR CAMPAIGN STRATEGY, FINANCE FOR SCENARIO MODELLING, OPERATIONS FOR SUPPLY CHAIN PLANNING.

“THE OUTPUTS ARE FLUENT, FAST, AND CONSISTENT. THE TEAMS TRUST THEM — AND HONESTLY, SO DO I. IT’S BECOMING DIFFICULT TO REMEMBER HOW WE WORKED BEFORE.”

— Chief Executive Officer, FMCG

Conclusion

Agency as Competitive Advantage

The integration of AI into strategic work is inevitable and, when properly governed, highly valuable.

But the competitive advantage in the AI era will not come from who adopts AI fastest or who engineers the best prompts, it will come from who preserves **cognitive agency**: the capacity to author reasoning, maintain verification discipline and orchestrate processes that leverage AI capabilities while protecting human judgment.

The AI Thinking Quadrants provide a diagnostic framework for this challenge and Cognitive Governance provides the mechanisms.

As intelligence becomes commoditised, the defining leadership capability is not finding answers—AI systems will do that with increasing competence. It is orchestrating the reasoning process itself: knowing which cognitive posture to adopt, when to switch, and how to preserve authorship even when machines participate in thinking.

A question that may distinguish high-performing organisations from those that drift is straightforward: **When you make consequential decisions using AI, who is actually doing the thinking?**

Cognitive Governance provides a valuable and timely framework for answering that very question.

Board Checklist

Cognitive Governance For AI-mediated Strategic Decisions

Before endorsing or circulating a strategic decision materially shaped by Generative AI, an organisation's board should require a minimum standard of cognitive "hygiene". The goal is not to slow down decision-making, but to prevent a common failure mode: mistaking fluent text for valid reasoning. The checklist below makes explicit who is doing the thinking, on what grounds, and under what verification discipline.

1. Decision rights (who can do what)

- Is it explicit where AI may generate options and where it may only assist?
- Who is the accountable human owner for framing, assumptions, and final choice?

2. Posture declaration (which Quadrant are we operating in?)

- Has the dominant posture been declared (Generative, Critical, Operating or Exploratory)?
- Is there evidence of posture drift (from "challenge my assumptions" to "tell me what to do")?

3. Switching discipline (mandatory sequence)

- Was a deliberate sequence followed: divergence → stress-test → convergence → execution planning?
- Was at least one Critical Posture pass required before commitment?

4. Transparent attribution (authorship clarity)

- What is AI-generated, what is human-generated, and what is human-validated?
- Are we presenting "our strategy" when the underlying narrative is primarily model-authored and merely curated by humans?

5. Evidence mapping (the justification chain)

- For each key claim: evidence, source, assumptions, verification status, residual uncertainty.
- Are the foundations primary (internal data/experts) or merely plausible (LLM narrative coherence)?

6. Assumption registry (points of failure)

- What are the five assumptions that, if false, would break the decision?
- Have we defined leading indicators or tests to validate them early?

7. Adversarial review (red-teaming)

- Do we have a credible, decision-relevant case against the chosen option?
- Was the challenge performed by an independent party, or by the same team that generated the output?

8. Verification capacity (human capability to verify)

- Who has the domain expertise to verify the high-risk components (market, legal, technical, regulatory)?
- Has at least one offline verification step been executed on the critical claims?

9. Cognitive audit trail (traceability of the process)

- Do we have a trace of postures adopted, sources used, iterations performed, and the final human rationale?
- Can the organisation reuse this trail for learning, not only for compliance?

10. Operational guardrails (before scaling)

- Are tool boundaries, data constraints, retention rules, and anti-shadow-AI policies explicit?
- Do we have a stop rule if signals of dependency or decision-quality degradation emerge?

11. Agency KPIs (monitoring erosion early)

- Are we tracking indicators such as posture distribution, frequency of deliberate switching, verification coverage, and near-misses?
- Is the organisation drifting toward convenience (Generative / Model-Led) at the expense of discipline (Critical / Human-Led verification)?

12. The non-negotiable question

- If this decision fails in 12–18 months, can we explain why we made it without defaulting to "the model recommended it"?

Cognitive Governance is not bureaucratic overhead. It is how Boards preserve the rarest capability in an era of commoditised intelligence: strategic agency and accountability for reasoning.

About the Authors

Paolo Taticchi is a Professor of Strategy and Sustainability and Deputy Director at UCL School of Management, where he co-directs the UCL Centre for Sustainable Business. A global expert in sustainability and strategy, he has trained thousands of Fortune 500 executives and developed business projects across five continents.

His widely cited research includes over 50 publications and books such as *How to Be Sustainable* (2025) and *Disruption* (2023). In 2025, his research on impact investing was highly commended by the Financial Times for making a real difference. A sought-after speaker and advisor, he has delivered 250+ talks attended by more than 100,000 people and serves on several international advisory boards.

As an entrepreneur, he co-founded four companies and led major educational initiatives. His accolades include *Poets & Quants* "Top 40 Under 40 Business Professors in the World" (2018), being mentioned by *Sole 24 Ore* as the most influential Italian under 40 (2021-23), and inclusion in the *Thinkers50 Radar List* (2025).

p.taticchi@ucl.ac.uk

www.paolotaticchi.com

Giovanni Giamminola is an independent researcher and author specialising in AI-driven business transformation. He serves as Founder and CEO of Systemic Zero, a strategic consulting firm focused on AI-First Enterprise transformation, and operates Systemic Zero Framework, a certification program on AI and leadership and management.

His client portfolio includes pharma, retail and industrial companies. He has authored ***The Augmented Manager***. He has taught executive education sessions at the **UCL School of Management** collaborating with UCL academics on programmes focused on AI and leadership and serves as Fellow UCL Institute for Sustainable Business.

giovanni.giamminola@systemiczero.com

www.systemiczero.com

UCL School of Management