Article

From Data Insight to Foresight: The Collaborative Logic and Boundary Constraints of Embedding Generative AI in Corporate Strategic Decision-Making

Wenxia Mao[1,*]

[1] Zhejiang Laboratory, Hangzhou, China
* Correspondence: Wenxia Mao, Zhejiang Laboratory, Hangzhou, China

Abstract: The entry of generative artificial intelligence into corporate strategy fundamentally transforms how organisations connect problem identification, evidence organisation, option development, and scenario exploration. Generative AI can convert dispersed materials into discussable strategic hypotheses and widen the search space for both opportunities and risks. However, its fluent output may also conceal hallucinated facts, misplaced causality, contaminated data, and ambiguous responsibility, thereby introducing new vulnerabilities into the strategic decision-making process. This paper examines how generative AI can support a progression from data insight to forward-looking judgement without being mistaken for an autonomous forecasting authority. Drawing on a conceptual analysis of human-AI interaction in strategic contexts, it argues that effective embedding requires a decision architecture built on verified evidence, competing scenarios, cross-examination between humans and models, and continuous feedback from realised outcomes. Access rights should vary according to the materiality, reversibility, and information sensitivity of each decision. Source attribution, human review, version records, dissent logs, and accountable sign-off must therefore operate within one auditable and transparent process. The study further proposes a governance framework that delineates the boundary conditions under which generative AI contributes to strategic foresight while preserving managerial accountability. Under these conditions, generative capability can become an enhanced strategic decision mechanism that remains open to verification, challenge, and correction rather than a substitute for managerial judgement. The findings offer actionable guidance for organisations seeking to integrate generative AI responsibly into their strategic planning and decision-making workflows.

Keywords: generative artificial intelligence; strategic decision-making; human-AI collaboration; foresight; algorithmic governance

Received: 19 June 2026
Revised: 29 July 2026
Accepted: 11 August 2026
Published: 16 August 2026

1. Introduction

The central difficulty in corporate strategic decision-making is often not a lack of information. It is the fragmentation of signals, disagreement over interpretation, and the limited visibility of long-term consequences. Market movements, customer responses, supply-chain fluctuations, regulatory adjustments, and other operational developments are recorded in separate systems and professional communities. Senior managers therefore receive facts that have already been filtered by departmental priorities. Generative AI provides a new cognitive interface because it can connect unstructured documents, summarise alternative narratives, and formulate questions across conventional information boundaries. Yet the same generative mechanism can complete missing links with plausible language when the underlying evidence is weak. The strategic value of the technology therefore depends less on textual fluency than on the institutional conditions under which its output is produced and challenged [1].

Strategic decisions combine uncertainty, cross-functional dependence, and strong commitment effects [2]. Machines are effective at search, calculation, and pattern detection, whereas managers remain responsible for contextual interpretation, balancing competing interests, and assessing the consequences of commitment. A productive relationship must therefore assign complementary tasks and create opportunities for mutual checking rather than pursue simple substitution. Research on human-AI collaboration indicates that stable managerial value emerges only when technology is incorporated into organisational processes and responsibility structures. For strategy, this means evaluating generative AI by whether it improves problem framing, expands option search, exposes assumptions, and preserves accountable judgement. The number of generated reports or the speed of drafting is not an adequate measure of strategic capability.

Treating generative AI merely as a reporting assistant leaves the decision process largely unchanged. Treating a coherent response as a reliable forecast is more dangerous because it can give unsupported assumptions the appearance of analytical authority [3]. Current research identifies organisational impact, risk governance, knowledge management, and collaborative optimisation as central issues in human-intelligence cooperation. The managerial task is consequently to design a system that moves from data insight to foresight while retaining the ability to detect and correct error. Such a system should make evidence visible, encourage competing interpretations, define the point at which human judgement becomes decisive, and learn from the difference between anticipated and realised outcomes.

2. How Generative AI Reconstructs the Information Basis of Strategic Decisions

Generative AI first changes how information enters the strategic agenda. Customer comments, patent documents, supplier communications, policy materials, and internal project records can be connected by topic rather than reviewed in isolation. The technology may reveal early indicators that have not yet appeared in financial metrics or formal market reports. Its role is not to replace statistical analysis but to transform heterogeneous material into analyzable questions, enabling data to serve not only as records of completed activity but also as evidence for forming strategic hypotheses [4]. This function is especially useful when a strategic issue spans several departments and no single data owner possesses a complete account.

Moving from insight to foresight requires a further shift in analytical framing. Instead of requesting the most likely outcome, managers should ask under what conditions a particular outcome could occur. Generative AI can organize baseline, stress, and plausible alternative scenarios around critical variables and specify possible triggers, transmission pathways, and leading indicators. These outputs are not forecasts in themselves. Their value lies in surfacing premises that were previously implicit and making them available for systematic comparison. Strategic discussion can then move away from competing positions and toward the quality of underlying assumptions, the credibility of supporting evidence, and the implications of potential misjudgment [5, 6].

Table 1 draws on a global AI survey conducted in 2025, which covered 1,491 respondents across 101 countries and weighted responses according to each country's contribution to global gross domestic product. The figures do not demonstrate that adoption necessarily produces value. They instead show the distance between broad use, process redesign, governance maturity, and measurable financial contribution [6].

Table 1. Key Indicators of Generative AI Adoption, Governance and Value Conversion in Enterprises

Observed indicator	Share (%)	Strategic implication
Regular use of generative AI in at least one business function	71	Adoption has entered operating processes
Fundamental redesign of at least some workflows	21	A substantial gap remains between use and process redesign
Chief executive responsible for AI governance	28	Strategic-level accountability remains limited

All generated content reviewed before use	27	Comprehensive human review remains insufficient
At least one negative consequence experienced	47	Risk controls need to be moved forward
Generative AI contributes at least 5% of enterprise EBIT	17	Enterprise-wide value has not been widely realised

Source: Compiled from McKinsey & Company's 2025 global AI survey. The survey was conducted from 16 to 31 July 2024, received 1,491 valid responses from 101 countries and weighted results by each country's contribution to global GDP.

Seventy-one per cent of respondents reported that their organisations regularly used generative AI in at least one business function, but only 21 per cent had fundamentally redesigned at least some workflows. Chief executive responsibility for AI governance and review of all generated content before use were reported by 28 and 27 per cent respectively. Meanwhile, 47 per cent had experienced at least one adverse consequence, and only 17 per cent attributed five per cent or more of earnings before interest and taxes to generative AI. Adoption therefore runs well ahead of process reconstruction and value realisation [1, 7]. The gap suggests that technical access becomes a strategic capability only when decision routines, governance structures, and accountability mechanisms are redesigned around it.

The managerial implication is that adoption indicators should be connected to an inventory of strategic decisions [3]. For each recurring decision, the enterprise should identify the information sources used, the stage at which a model participates, and the evidence required before commitment. This inventory reveals where generative AI merely accelerates document production and where it changes the quality or timing of judgement. It also exposes decisions that are technically suitable but institutionally unprepared because no designated owner can verify evidence or resolve disagreement. Linking use cases to decision rights creates a more credible basis for investment than estimating value from the number of active users.

3. Collaborative Logic from Data Insight to Foresight

3.1. Constraining Generation with Verifiable Evidence

An enterprise should define the admissible evidence before selecting a model or writing prompts. Operating data, policies, project reviews and verified external materials can be organised in a tiered knowledge base. Each item should carry its source, date, responsible department and scope of validity. Model output should distinguish facts, inferences and recommendations [3, 8]. Facts require traceable support, inferences require stated conditions, and recommendations require explicit boundaries. Missing data should be identified rather than silently completed. Where departmental records differ, both versions should remain visible until an accountable owner resolves the discrepancy. These practices prevent generic model knowledge from displacing the organisation's actual circumstances.

The knowledge base also requires update frequencies, expiry rules and a procedure for handling discrepancies. For strategically important conclusions, the system should present supporting evidence, counter-evidence and unresolved uncertainty together [4]. This arrangement keeps the model in the role of an evidence organiser rather than an authority on corporate facts. Managers can then separate what is known from what is inferred and decide whether the evidential gap can be tolerated. Human-AI complementarity arises through contextual understanding and bias correction, not from accepting machine output at face value. Evidence governance is therefore the first control point in an enhanced strategic decision process.

Evidence governance also has a temporal dimension [9]. A source that was reliable during one market cycle may become misleading after shifts in policy frameworks, pricing dynamics or customer behaviour. Retrieval systems should therefore preserve effective dates and prevent outdated documents from receiving the same weight as current records.

Important prompts should state the time horizon of the decision and flag evidence produced under different conditions. Where strategy depends on rapidly changing external information, an owner should confirm that the knowledge base was refreshed before review. Temporal control reduces the risk that a model produces a coherent answer from facts that are individually accurate but no longer aligned with current conditions.

3.2. Expanding Option Search through Competing Scenarios

A single narrative can dominate foresight even when it rests on fragile assumptions. Generative AI can use the same evidence set to construct interpretations from the perspectives of customers, competitors, suppliers, regulators, and internal operating units. It can then integrate critical uncertainties into contrasting scenarios and compare the revenue potential, risk exposure, capability requirements, and exit costs associated with alternative actions. The objective is not to produce a longer report, but to prevent premature convergence and ensure each option withstands multiple causal interpretations. Scenarios are especially valuable when strategic disagreement arises from differing assumptions rather than differing preferences, as they clarify the root cause of disagreement.

Scenario analysis becomes actionable only when broad narratives are translated into concrete decision points. Each scenario should specify observable triggers, leading indicators, required resources, a timeframe within which decisions remain reversible, and clear criteria for pausing or halting resource commitments. Analytical models may reconfigure variables and action combinations, but managers must assess whether the underlying conditions are realistic and whether the cost of commitment is justified. When a leading indicator shifts, the organization should first adjust the relative likelihood of each scenario and then determine whether to accelerate, delay, revise, or terminate related resource allocations. In this way, foresight serves as preparation for contingent action rather than a claim to predictive certainty.

Scenario breadth must be balanced with practical decision utility. Generating numerous narratives can create an illusion of rigor while diluting managerial focus. The enterprise should retain a limited set of scenarios that differ meaningfully in causal logic—not merely in tone or degree of optimism or pessimism. Each retained scenario must identify at least one decision that would change if that scenario gained greater plausibility. If two scenarios lead to identical actions, they may not require separate governance structures. This discipline ensures generative exploration remains tightly linked to resource allocation and avoids turning scenario development into a purely presentational exercise.

3.3. Calibrating Strategic Judgement through Cross-Examination

Model output must undergo a structured challenge process prior to informing any decision. Evaluations generated by a single large language model are sensitive to prompt formulation, assigned roles, and information sequencing. Aggregating assessments across multiple models and functional perspectives yields judgements that better approximate expert evaluation. Enterprises may therefore decouple generation, critical appraisal, and managerial oversight: one model proposes an option, another identifies potential weaknesses, while domain specialists—such as those in business strategy, finance, technology, legal affairs, and risk management—examine assumptions within their respective areas of responsibility. Disagreement in this process is not indicative of failure; rather, it signals the need for deeper scrutiny.

For high-stakes decisions, prompts, evidence sets, and model versions must be formally locked before review commences. Supporting and opposing teams must formulate positions independently, and meeting records must preserve substantive objections that were ultimately not adopted. All challenges must explicitly reference empirical evidence, underlying assumptions, or defined decision boundaries—not opaque technical scores. Human reviewers are responsible for verifying the integrity of the evidence chain, coherence of causal reasoning, scope of authority, and potential distribution of adverse impacts. The final signatory must articulate the rationale clearly in original language, without paraphrasing or reproducing the model's output—ensuring accountability remains with human decision-makers.

3.4. Converting Outcome Feedback into Organisational Learning

Following action, the enterprise should document the initial hypothesis, triggering scenarios, actual outcomes, and root causes of deviation. It must differentiate among information gaps, model inaccuracies, managerial judgments, and execution variances. A review should not assess decision quality solely on success or failure, as sound decisions may yield adverse outcomes under uncertainty. The more meaningful evaluation concerns whether the decision was reasonable given the evidence available at the time and whether the system identified errors early enough to constrain correction costs. Feedback constitutes organisational learning only when it leads to adjustments in subsequent evidence requirements, scenario weighting, system prompts, review responsibilities, or authority protocols.

4. Boundary Constraints on Generative AI in Strategic Decision-Making

4.1. Boundaries of Fact and Causality

Generative models organize language through learned statistical associations. They lack built-in access to authoritative corporate data, and textual correlations do not imply causal relationships. In strategic decisions—such as market entry, major capital allocation, or technology roadmap planning—a model may synthesize evidence and generate plausible explanations, but it cannot independently verify actual demand levels or anticipate competitor behavior. Quantitative assertions must be verified against primary source records. External claims require explicit attribution to their origin, while causal assertions necessitate independent validation via analytical reasoning, empirical testing, or rigorous due diligence. A coherent narrative does not guarantee factual reliability. The greater the apparent clarity of an output, the more essential it becomes to distinguish between directly observed information, logically inferred conclusions, and unresolved uncertainties.

4.2. Boundaries of Data and Intellectual Property

Strategic materials may contain trade secrets, personal information, and undisclosed financial data. Transmitting such information to public models risks premature exposure prior to any formal decision-making. Enterprises must classify information according to sensitivity levels and strictly govern permissible models, user accounts, and access methods for each category. Sensitive fields may require data masking, principle-of-least-privilege access controls, and comprehensive audit logging. Supplier agreements must explicitly prohibit the use of customer inputs for model training purposes. Regulatory frameworks in China assign responsibility for data security, personal information protection, intellectual property rights, and content authenticity to providers of generative AI services. Internal policies should operationalize these requirements through defined input inventories, prohibited usage scenarios, and mandatory review checkpoints.

4.3. Boundaries of Authority and Responsibility

Generative AI can propose options but cannot assume responsibility for financial loss, employment outcomes, or regulatory compliance obligations. Decision-making authority should be calibrated according to the materiality, reversibility, and sensitivity of information involved. For low-risk, reversible operational experiments, model-generated recommendations may be implemented following business validation and within a pre-approved budget [10, 11]. In contrast, decisions involving long-term capital allocation, core technology development, major strategic acquisitions, or organizational restructuring require formal review by designated personnel with explicit accountability. The model's involvement should diminish as decision impact and irreversibility increase. Managers remain fully accountable for assumptions they endorse and actions they authorize; citing model output does not relieve them of this responsibility.

4.4. Boundaries of Organisational Cognition

Heavy reliance on similar models can cause departmental responses to converge and suppress authentic disagreement [12]. Strategic judgement must incorporate front-line experience, direct customer interaction, and domain-specific expertise. Proficiency in using models does not replace deep industry understanding. Different teams may independently analyse the same issue before comparing conclusions, thereby mitigating the risk that shared prompts generate artificial consensus. Guidelines for labelling AI-generated or synthetic content require clear marking—whether explicit or implicit—and comprehensive provenance documentation. Internal strategic materials should similarly disclose whether content was generated or substantially revised by a model, specify the model version used, and record the responsible processor and review status.

The cognitive boundary also encompasses the risk of gradual deskilling. Repeated delegation of evidence synthesis and option framing to automated systems may erode managers' ability to detect omissions or formulate independent alternatives. Review teams should periodically undertake selected decisions without generative assistance and compare the resulting reasoning process with AI-supported outputs. Rotating between model-assisted and independent analysis helps sustain professional judgement. Training programmes should emphasise reconstructing answers directly from source material, not solely refining prompting techniques. The aim is to expand cognitive capacity while preserving the human capability to critically evaluate and challenge system outputs.

5. Building an Auditable, AI-Enhanced Strategic Decision Mechanism

Organisational embedding should begin with the decision task rather than the technical feature [13]. Tasks may be grouped by impact and reversibility into information organisation, option exploration, operating experiments, and major commitments. Models may participate deeply in the first two categories if evidence is attached. Operating experiments require a budget, time limit, and stopping condition. Major commitments require model assistance to remain subordinate to multi-person review and senior sign-off. This tiered structure prevents convenience from expanding model authority beyond an acceptable boundary.

Major matters should produce a standard decision-evidence package. It should include problem definition, data scope, key assumptions, competing scenarios, objections, leading indicators, and exit rules [14]. Business units should be accountable for facts and implementation conditions. Data teams should maintain data quality and model records. Legal and risk functions should review regulatory and operational requirements, while a decision committee remains responsible for the final trade-off. A change in the model version, evidence scope, or a critical assumption should trigger renewed evaluation rather than be treated as a routine technical update.

Governance can combine central standards with distributed decision responsibility. The enterprise level should define model admission, data classification, prompt templates, logging formats, and incident procedures [15, 16]. Business units should retain responsibility for problem definition, scenario selection, and resource commitment. A governance committee should review model drift, authority use, and recurring risk events, then update common rules where necessary. Supplier oversight should cover data processing, service continuity, and model-update terms. User training should emphasise evidence awareness, causal reasoning, and responsibility, not merely efficient prompting.

Auditability further requires an incident-learning loop. A material hallucination, unauthorised disclosure, or misleading recommendation should be recorded with the model version, prompt, retrieved material, reviewer, and decision impact. The response should address both the immediate case and the control weakness that allowed it to pass. Possible actions include changing access rules, improving source quality, revising prompt templates, or narrowing the approved task. Repeated incidents should trigger reconsideration of the model or supplier rather than additional warnings to users. This process turns risk events into evidence for improving the decision architecture.

Evaluation should shift from use volume to decision quality. Evidence traceability, hypothesis validation, preservation of dissent, warning lead time, correction cycles, and risk-event rates provide more meaningful indicators than prompt counts. Table 1 shows that only 27 per cent of respondents reviewed all generated content, whereas 47 per cent had encountered negative consequences. The 17 per cent reporting a material earnings contribution must also be interpreted alongside investment, process redesign, and risk exposure. Short-cycle experiments should emphasise early warning and correction cost, while long-term commitments require monitoring of assumption failure and timely resource adjustment.

Implementation can begin with issues that have clear evidence, observable outcomes, and reversible action [17]. A human-only baseline should be retained so that information coverage, option diversity, decision time, and subsequent deviation can be compared. Before a pilot is scaled, managers should determine whether improvement came from the model, better data, or process redesign. Maturity is not demonstrated by the number of meetings attended by a model. It is demonstrated when the organisation can explain the evidence used, the parties who challenged it, the person who decided, and the mechanism by which error will be corrected.

The same discipline should govern retirement as well as scaling. A model may cease to be suitable because its supplier changes, its knowledge boundary no longer matches the decision context, or its correction burden exceeds the value obtained. Periodic renewal should therefore compare the application with a human-only process and with available alternatives [16]. Withdrawal criteria protect the enterprise from treating prior investment as proof of continuing usefulness. They also keep strategic foresight responsive to changes in technology and business conditions rather than locking decision practice to an early adoption choice.

6. Conclusion

Generative AI can advance corporate strategic decision-making by connecting evidence, generating scenarios, and sustaining structured challenge. It cannot eliminate uncertainty or substitute for managerial responsibility. Progressing from data insight to foresight requires verified evidence to constrain generation, competing scenarios to surface assumptions, cross-examination to mitigate bias, and outcome feedback to update organisational knowledge. Tiered authority, human review, provenance records, version control, dissent logs, and accountable sign-off must be developed in concert. When embedded in the decision process, generative capability becomes a strategic capability that is verifiable, contestable, and correctable. Without these safeguards, faster production of persuasive text may heighten rather than reduce strategic risk.

References

1. H. Li and F. Tian, "Advancing decision-making through AI-human collaboration: A systematic review and conceptual framework," *Group Decision and Negotiation*, vol. 35, no. 2, p. 26, 2026.
2. C. Anthony, B. A. Bechky, and A. L. Fayard, "'Collaborating' with AI: Taking a system view to explore the future of work," *Organization Science*, vol. 34, no. 5, pp. 1672–1694, 2023.
3. B. Lou, T. Lu, T. S. Raghu, and Y. Zhang, "Unraveling human-AI teaming: A review and outlook," arXiv preprint arXiv:2504.05755, 2025.
4. A. Singla, A. Sukharevsky, L. Yee, M. Chui, and B. Hall, *The State of AI: How Organizations Are Rewiring to Capture Value*. McKinsey & Company, 2025, p. 12.
5. A. Trunk, H. Birkel, and E. Hartmann, "On the current state of combining human and artificial intelligence for strategic organizational decision making," *Business Research*, vol. 13, no. 3, pp. 875–919, 2020.
6. A. R. Doshi, J. J. Bell, E. Mirzayev, and B. S. Vanneste, "Generative artificial intelligence and evaluating strategic decisions," *Strategic Management Journal*, vol. 46, no. 3, pp. 583–610, 2025.
7. O. López-Solís et al., "Effect of generative artificial intelligence on strategic decision-making in entrepreneurial business initiatives: A systematic literature review," *Administrative Sciences*, vol. 15, no. 2, p. 66, 2025.
8. M. Albashrawi, "Generative AI for decision-making: A multidisciplinary perspective," *Journal of Innovation & Knowledge*, vol. 10, no. 4, p. 100751, 2025.
9. P. Korzynski et al., "Generative artificial intelligence as a new context for management theories: analysis of ChatGPT," *Central European Management Journal*, vol. 31, no. 1, pp. 3–13, 2023.

10. P. Ravichandran Sr, J. R. Machireddy Sr, and S. K. Rachakatla, "Harnessing generative AI for automated data analytics in business intelligence and decision-making," *Hong Kong Journal of AI and Medicine*, vol. 4, no. 1, pp. 122–145, 2024.
11. S. V. Khara et al., "Generative Artificial Intelligence for Intelligent Decision Support in Organizational Management: A Data-Driven Framework," in *Proc. 2026 3rd Int. Conf. Res. Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE)*, Apr. 2026, pp. 1–5.
12. M. A. Toha, P. A. Khan, and M. S. Uddin, "Theorizing Generative AI's Role in Strategic Decision-Making: From Automation to Augmentation," in *INFORMS Int. Conf. Service Science*, Jul. 2025, pp. 114–125.
13. B. Bapat et al., "Leveraging Generative AI for Strategic Decision-Making: Enhancing Managerial Effectiveness in Dynamic Business Environment," in *Proc. 2025 IEEE Int. Conf. Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI)*, Mar. 2025, vol. 3, pp. 1–6.
14. N. Drašković, M. Rešetar, and V. Vejzagić, "Managerial Acceptance of Generative Artificial Intelligence for Strategic Decision Making: A Review of Evidence and Frameworks," in *Proc. 2026 49th MIPRO ICT and Electronics Convention (MIPRO)*, May 2026, pp. 1166–1171.
15. N. Schulte and D. K. Kanbach, "Rethinking organizational decision-making: The emerging roles and tasks of generative artificial intelligence," *Management Review Quarterly*, pp. 1–68, 2026.
16. M. Ayyoobi, "The Role of Artificial Intelligence in Strategic Decision Making and Business Innovation," *Int. J. Adv. Stud. Hum. Soc. Sci.*, vol. 14, no. 2, pp. 144–156, 2025.
17. C. P. Gupta and V. R. Kumar, "Enhancing Decision Intelligence through Generative AI: A Framework for Data-Driven Strategic Planning," in *Proc. 2026 Int. Conf. Commun., Comput. and Emerging Technol. (IC3ET)*, Feb. 2026, pp. 596–601.