

The invisible hand of generative AI: quasi-agency and knowledge transformation in data-driven decision-making

Matteo Cristofaro, Alexis Bañón-Gomis and Pier Luigi Giardino 

Abstract

Purpose – This study aims to examine how generative artificial intelligence (GenAI) reshapes knowledge management (KM) in data-driven decision-making (DDDM), reconfigures principal-agent relations and creates new challenges for epistemic governance. It focuses on how GenAI influences the construction, synthesis, justification and retention of decision-relevant knowledge across the decision cycle.

Design/methodology/approach – This study adopts an abductive qualitative design, drawing on data from 120 managers at Italian small and medium-sized enterprises actively using GenAI in recurring DDDM activities. Drawing on ten open-ended questions aligned with the decision cycle, the analysis applies reflexive thematic analysis, supported by Gioia-informed coding procedures, to identify recurring phase-level governance patterns and to develop a phase-sensitive framework.

Findings – This study develops a phase-sensitive model of GenAI quasi-agency showing how generative artificial intelligence reshapes organizational knowledge throughout the decision cycle. The model explains that GenAI transforms knowledge through three recurring mechanisms – knowledge coupling, knowledge decoupling and knowledge hiding – which operate with different intensity across decision phases and redefine the governance requirements for traceability, reconstructability, contestability and human justificatory ownership. Consequently, agency risks increasingly originate from the transformation of hidden knowledge rather than from hidden human action.

Originality/value – This study introduces the concept of GenAI quasi-agency, defined as the capacity of GenAI systems to shape framing, evidence construction, prioritization, justification and organizational memory without possessing formal authority, intentionality or accountability. It contributes to KM by conceptualizing GenAI-supported decision-making as a problem of epistemic governance centered on preserving traceability, reconstructability, contestability and human justificatory ownership. The study also extends principal-agent theory by showing how agency risks increasingly arise from the transformation of hidden knowledge rather than from hidden action alone and develops a phase-sensitive framework that explains how governance requirements vary across the DDDM cycle.

Keywords Generative AI, Knowledge management, Principal-agent theory, Data-driven decision-making, SMEs

Paper type Research paper

Matteo Cristofaro is based at the Department of Management and Law, University of Rome Tor Vergata, Rome, Italy. Alexis Bañón-Gomis is based at the Department of Business Organization, Universitat Politècnica de València, Valencia, Spain. Pier Luigi Giardino is based at the Department of Business Administration, Pontificia Universidad Javeriana, Bogotá, Colombia.

Received 23 March 2026
Revised 25 May 2026
24 June 2026
Accepted 3 August 2026

© Matteo Cristofaro, Alexis Bañón-Gomis and Pier Luigi Giardino. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at <http://creativecommons.org/licenses/by/4.0/>

1. Introduction

Generative artificial intelligence (GenAI) refers to artificial intelligence (AI) systems that produce textual, visual or analytical outputs through probabilistic generation based on large-scale training data and is becoming embedded in organizational knowledge and decision infrastructures. In small and medium-sized enterprises (SMEs), managers increasingly use GenAI to retrieve evidence, synthesize information, generate alternatives, compare options, interpret performance and support learning across decisions (Cristofaro and Giardino, 2025; Hernández-Tamurejo *et al.*, 2025). These activities place GenAI within data-driven decision-making (DDDM), a structured process in which data are collected,

interpreted and transformed into evidence to support problem definition, the evaluation of alternatives, justification, implementation and the revision of organizational choices (Kampoowale, 2025). Recent knowledge management (KM) research further suggests that GenAI expands organizations' ability to access, recombine and mobilize distributed knowledge across the processes of acquisition, retrieval, sharing and application (Leoni *et al.*, 2024; Zhang *et al.*, 2025; Cristofaro and Bañón-Gomis, 2026).

Yet, the central issue is not simply faster information processing or greater analytical efficiency. GenAI not only processes information; it also transforms it into organizationally legitimate knowledge for action. Generated outputs summarize, rank, compress, translate and reframe evidence before managers use them to define problems, evaluate alternatives, justify decisions, interpret performance or update organizational memory. A generated summary may detach claims from their evidentiary context and a generated ranking may confer disproportionate legitimacy on one ordering of alternatives while compressing uncertainty, disagreement or contextual nuance (Kellogg *et al.*, 2020; Leoni *et al.*, 2024; Zhang *et al.*, 2025). The governance challenge therefore shifts from information availability toward epistemic governance, namely, the governance of processes that preserve traceability, reconstructability, contestability and human accountability as information is transformed into legitimate knowledge for action (Nonaka, 1994; Grant, 1996).

This creates an important theoretical problem for both KM and principal-agent theory. KM research has traditionally distinguished information from actionable knowledge, emphasizing interpretation, contextualization, validation and organizational learning as central knowledge processes (Argote and Miron-Spektor, 2011). Yet GenAI complicates these distinctions because fluent outputs can appear authoritative even when their assumptions, uncertainties or inferential pathways are difficult to reconstruct. At the same time, classic agency theory focuses primarily on hidden action, information asymmetry, monitoring and incentive alignment (Jensen and Meckling, 1976; Eisenhardt, 1989; Bendickson *et al.*, 2016). In GenAI-supported decision-making, however, agency concerns extend to the transformation of hidden knowledge, where actors rely on coherent summaries, rankings or rationales without understanding how evidence was selected, reframed or legitimized. We therefore ask:

Q. How does GenAI shape the transformation of decision-relevant knowledge and what implications does this have for principal-agent relations?

Italian SMEs that actively use GenAI provide a revealing setting because leaner governance and less layered decision processes make the movement from evidence to rationale comparatively visible. The study draws on structured accounts from 120 managers, analyzed through abductive reflexive thematic analysis with Gioia-informed coding, enabling a transparent movement from accounts to themes and a phase-sensitive framework (Braun and Clarke, 2006, 2019; Gioia *et al.*, 2013).

Our analysis develops a phase-sensitive framework of GenAI quasi-agency that explains how knowledge is transformed across the DDDM cycle. The framework identifies three recurrent knowledge dynamics. GenAI supports knowledge coupling by integrating evidence, assumptions, rationales, actors and prior learning into shared decision infrastructures. It produces knowledge decoupling when summaries, rankings or rationales detach from their evidentiary basis, while prompting governance safeguards that separate AI-assisted synthesis from human authorization. It contributes to knowledge hiding when retrieval, summarization, ranking or framing selectively reduces the visibility of caveats, dissent, uncertainty, alternatives or counterevidence (Connelly *et al.*, 2012; Kellogg *et al.*, 2020; Arias-Pérez and Vélez-Jaramillo, 2022). The framework further shows that these three dynamics vary in salience across the decision cycle, shaping the epistemic conditions under which organizational knowledge remains traceable, reconstructable, contestable and subject to human justificatory ownership. We conceptualize GenAI quasi-agency as the

capacity to shape framing, evidence construction, prioritization, justification, performance interpretation and organizational memory without formal authority, intentionality or accountability (Ågerfalk, 2020; Dattathrani and De, 2023; Humberd and Latham, 2025).

This study makes three contributions. First, it contributes to KM research by conceptualizing GenAI-supported decision-making as not only a matter of information retrieval but also an epistemic governance problem. Specifically, it explains how knowledge coupling, decoupling and hiding redirect attention toward the conditions necessary for AI-supported decisions to remain traceable, reconstructable, contestable and accountable to human judgment. Second, it extends principal-agent theory by showing that agency risks in GenAI-supported DDDM increasingly arise from hidden knowledge transformation rather than from hidden action or limited effort observability alone (Jensen and Meckling, 1976; Eisenhardt, 1989; Bendickson *et al.*, 2016). Third, it contributes to DDDM research with a phase-sensitive framework showing how GenAI's agency-relevant influence varies across framing, evidence construction, prioritization, implementation, measurement and organizational memory (Kampoowale, 2025). Taken together, these contributions reposition GenAI from a productivity aid to a governance-relevant participant in organizational knowledge transformation.

2. Theoretical background

2.1 Agency theory and (artificial) agency in organizations

Agency theory begins with the separation of ownership and control. Delegation of decision rights generates monitoring, bonding and residual costs under divergent interests, information asymmetries and uncertainty (Jensen and Meckling, 1976). Contract design balances outcome-based incentives and behavior-based monitoring according to measurability, risk sharing and uncertainty (Eisenhardt, 1989). In AI-mediated settings, agency problems extend beyond the principal-agent dyad because algorithmic systems shape the information, categories and representations through which actors define, justify and evaluate decisions (Stelmaszak *et al.*, 2026).

Agency theory's assumptions about observability, measurability and accountability are strained in AI-mediated decision-making. Traditional problems centered on unobserved actions or intentions, but in AI-mediated settings decision-relevant information may already have been selected, summarized, ranked, transformed or stabilized through algorithmic systems before reaching managerial judgment. Agency governance therefore concerns the conditions under which information becomes observable, comparable, auditable and usable for action.

Research on digital and artificial agency provides a foundation for understanding this shift. Digital agency emphasizes enactment through code, data and sociotechnical arrangements, while artificial agency emphasizes autonomy, adaptivity and interaction in human-machine environments (Ågerfalk, 2020; Dattathrani and De, 2023). Agency in digital contexts cannot be reduced to human intention alone. GenAI's governance relevance lies less in decision authority than in the production and transformation of decision-relevant outputs. GenAI systems may summarize evidence, generate alternatives, rank priorities and retrieve prior organizational knowledge without owning the decision or bearing accountability.

Algorithmic systems also reshape what becomes observable, contestable and sanctionable by directing action, evaluating performance and structuring categories (Kellogg *et al.*, 2020). Opacity may derive from strategic secrecy, limited user understanding or technical complexity. When transparency is infeasible, governance depends on evaluability, auditability, traceability and credible assessment of how outputs enter organizational decisions rather than only on inspection of model mechanics.

GenAI also affects incentive alignment and risk sharing. Alignment extends to hybrid human-machine arrangements involving reward functions, override authority, logging, explainability and escalation (Humberd and Latham, 2025; Stelmaszak *et al.*, 2026). Prediction, synthesis and classification may make some outcomes more measurable, but opaque or drift-prone systems shift risks to actors who cannot verify how outputs are generated (Eisenhardt, 1989; Dattathrani and De, 2023).

Trust adds a further governance challenge. Attributing agency to GenAI may shift trust toward the system itself and intensify betrayal aversion when expectations fail (Vanneste and Puranam, 2025). Under opacity, audit access, certification and accountable intermediaries may stabilize trust without a direct understanding of model mechanics. Accountability, therefore, depends not only on who authorizes a decision but also on how generated outputs are interpreted, challenged and legitimized within organizational routines (Arnaout *et al.*, 2026).

These developments create a conceptual problem for principal-agent theory. Classic agency theory treats information asymmetry as hidden action, hidden information or limited observability of effort. In GenAI-supported decision-making, asymmetry may also arise from hidden knowledge transformation, where managers receive coherent summaries or rationales without being able to reconstruct how evidence was selected, weighted or reframed. Governance therefore extends from monitoring behavior and outcomes to governing how evidence becomes rationale, memory and legitimate grounds for action.

2.2 Knowledge management under generative AI

KM explains how organizations create, integrate, retain, transfer and apply knowledge for coordinated action under specialization and uncertainty. It draws on the knowledge-based view of the firm (Kogut and Zander, 1992; Grant, 1996) and on organizational learning research showing how knowledge becomes embedded in routines and memory systems (Argote and Miron-Spektor, 2011). KM distinguishes explicit from tacit knowledge, where value arises only through interpretation and enactment in practice (Nonaka, 1994). A longstanding concern is the conflation of storage with knowledge, reducing KM to managing artifacts rather than governing knowledge-in-use (Mårtensson, 2000).

GenAI intensifies this problem by affecting not only retrieval and storage but also the production, recombination, summarization and circulation of knowledge claims. Fluent outputs encourage treating synthesized content as validated knowledge rather than provisional interpretation. The KM challenge thus shifts from information access toward governing the transformations through which evidence is stabilized as decision-relevant knowledge (Cristofaro and Bañón-Gomis, 2026). AI's effects on resilience, innovation and performance are mediated through KM processes of acquisition, sharing, integration and application, though GenAI introduces risks such as hallucinations, overload and bias (Li *et al.*, 2025; Zhang *et al.*, 2025). Outcomes depend less on output fluency than on whether generated claims remain contextualized, validated and reconstructable.

Tacit and experiential knowledge is situated and context-dependent and cannot be fully captured by textual or statistical representations (Nonaka, 1994). GenAI supports externalization of codified knowledge but cannot transfer tacit knowledge through data alone and AI remains limited in reproducing empathy, creativity, improvisation and contextual judgment (Trio *et al.*, 2026).

Users may over-rely on fluent outputs or disengage from claims they cannot verify when source reliability and assumptions are opaque (Li and Yuan, 2026). Transparency, explainability and traceability become critical when systems compress reasoning into decision-ready outputs (Leoni *et al.*, 2024). GenAI may strengthen sharing when seen as augmenting expertise but discourage contribution when employees view AI as threatening

their knowledge (Chen *et al.*, 2026). KM under GenAI extends beyond distribution to the governance of interpretive authority, epistemic ownership and accountability.

Prior research separates automated, augmented and supported AI involvement by complexity and authority (Leoni *et al.*, 2024). Cristofaro and Bañón-Gomis (2026) show that AI-assisted decisions depend on preserving framing sovereignty, conditional delegation and feedback routines. GenAI therefore shifts KM from managing knowledge availability toward governing traceability, contestability, validation and delegation. The central challenge is how information becomes legitimate grounds for action and who owns the resulting rationale.

3. Methodology

3.1 Research design and data collection

This study uses an abductive qualitative design to examine how managers describe GenAI's role in KM and DDDM within Italian SMEs. We combine reflexive thematic analysis with Gioia-informed coding to move transparently from written accounts to theoretical themes (Braun and Clarke, 2006, 2019; Gioia *et al.*, 2013). Principal-agent theory and KM served as sensitizing, not fixed, coding categories. The DDDM cycle, rooted in information-processing theory and evidence-based management, informed questionnaire design and ensured comparable coverage across decision activities (Brynjolfsson *et al.*, 2011).

Data were collected via a structured online questionnaire administered between January and February 2025 to managers of Italian SMEs. Respondents received information regarding the study's purpose, voluntary participation, confidentiality, anonymization and informed consent. Participants were recruited through university networks, professional associations and the authors' contacts and the study complied with institutional ethical guidelines for research involving human participants.

We used purposive sampling to identify managers whose firms had incorporated GenAI into recurring decision workflows rather than into experimentation or peripheral administrative tasks, to obtain analytically rich accounts of AI-mediated managerial practices. Eligibility required that respondents held a managerial or executive role, had direct involvement in data-supported decision-making, could describe concrete uses of GenAI in organizational decision activities and represented an SME according to the European Commission definition. Active GenAI use was verified through screening questions about recurring decision workflows and the DDDM activities in which GenAI was applied.

A total of 123 SMEs initially met the eligibility criteria. Three firms did not return the questionnaire, leaving a final analytic sample of 120 managers from as many Italian SMEs, with one respondent per firm. Further details on sampling are reported in Supplementary Material SM1.

Supplementary Material SM2 reports descriptive statistics. Respondents were largely mid-career decision-makers, with most aged 30–50 ($n = 119$) and a gender distribution of 70% men and 30% women. They occupied senior organizational roles, with 75% top managers and 25% executives. The sample was industry-diverse, with 93 respondents reporting advanced GenAI experience and 27 reporting intermediate experience. This profile should be considered when interpreting the findings, as the accounts may reflect more sophisticated GenAI use than would be expected in firms with lower AI familiarity.

The questionnaire consisted of two sections. The first gathered organizational and respondent characteristics. The second included ten open-ended questions aligned with the seven DDDM steps (Cristofaro *et al.*, 2025); the additional questions unpacked some activities into more fine-grained practices related to interpretation, evaluation, prioritization and results assessment. Questions were neutrally worded to avoid embedding theoretical constructs into respondents' accounts. Supplementary Material SM3 reports the number of responses, mean and median response length, response-length range, degree of

elaboration and evidence-depth indicators for each item; Supplementary Material SM4 reports full anonymized response examples.

3.2 Data analysis

Data analysis followed abductive reflexive thematic analysis with Gioia-informed coding (Braun and Clarke, 2006, 2019; Gioia *et al.*, 2013). We do not claim to develop inductive process theory in the strict Gioia sense; Gioia-informed procedures served as transparency devices for moving from informant expressions to theoretically interpreted themes.

Each response was read and coded in full. When a response contained more than one relevant idea, multiple meaning units were retained rather than reducing the response to a single fragment. Meaning units were coded when they referred to GenAI use, KM, evidence construction, accountability, authority or one of the DDDM activities, allowing responses to contain multiple interpretive codes where appropriate.

The research team first read all responses several times and wrote initial notes on recurring references to decision authority, evidence use, accountability, implementation, measurement and organizational learning (Braun and Clarke, 2006, 2019). During coding, team members independently reviewed the responses. Independent coding generated alternative interpretations rather than intercoder reliability statistics, consistent with reflexive thematic analysis, which treats coding as interpretive rather than mechanical (Braun and Clarke, 2019). Codes stayed close to respondents' language where possible (Gioia *et al.*, 2013) and captured AI-assisted framing, evidence aggregation, traceability checking, scenario generation, ranking support, human-authored justification, escalation rules, override rights, source-of-truth separation and after-action learning.

Candidate themes grouped related first-order codes into broader patterns based on recurring meanings across responses, firms and DDDM activities rather than frequency alone. Analytic memos documented coding choices, tensions and deviant cases, while constant comparison refined codes and distinctions (Glaser and Strauss, 1967; Miles *et al.*, 2014). Themes were reviewed against coded extracts and full responses and were split, merged, renamed or refined when needed. The review included negative cases, such as accounts where GenAI had limited influence, was bypassed or was used mainly for productivity support, strengthening credibility through negative-case analysis (Lincoln and Guba, 1985; Miles *et al.*, 2014).

The analysis remained abductive: principal-agent theory and KM-informed interpretation guided interpretation, but did not define predetermined coding categories (Eisenhardt, 1989). This process yielded sixteen second-order themes and four aggregate dimensions, documented in a formal data structure that supports interpretive traceability (Gioia *et al.*, 2013).

We organized the themes into a phase-sensitive framework structured around the DDDM cycle. Because the ten open-ended questions were aligned with DDDM activities, responses were initially linked to the corresponding activity and where responses crossed phases, relevant meaning units were coded to the appropriate phase. Phase allocation was guided by temporal orientation and functional focus: problem definition, data preparation, interpretation and alternative generation, selection, implementation, measurement or review. Two researchers reviewed phase allocation independently before the team resolved disagreements by returning to full responses. The framework is therefore an interpretive reconstruction of reported practices, not direct evidence of causal governance change (Langley, 1999).

4. Results

4.1 Thematic structure of the findings


The questionnaire generated 1,200 written managerial accounts across the ten open-ended questions. Consistent with the abductive reflexive thematic approach described above,

each response was read in full. When respondents addressed more than one relevant issue within the same response, multiple meaning units were retained rather than reducing the response to a single analytically salient clause. This allowed the analysis to preserve the multidimensional character of the accounts, especially where respondents simultaneously discussed GenAI use, evidence construction, authority, accountability and DDDM activities (Braun and Clarke, 2006, 2019).

The analysis produced 160 first-order concepts, which were formulated as closely as possible to respondents' language. These concepts were then consolidated into 16 s-order interpretive themes through iterative comparison across responses, firms and DDDM activities. The themes were organized into four aggregate dimensions, which represent the highest level of conceptual integration in the data structure (see Figure 1). This data structure is presented to make the movement from respondent expressions to interpretive themes transparent, not to claim a strictly inductive Gioia process-theory design.

This coding architecture shows how the analysis moved from written managerial accounts to theoretically interpreted themes while preserving traceability between respondent expressions, interpretive themes and the phase-sensitive framework developed below (Gioia et al., 2013).

Three recurring patterns became central to the analysis. Some accounts described practices that kept decision-relevant knowledge connected to evidence, assumptions, criteria, actors and prior learning; others described risks or safeguards involving the detachment of generated outputs from sources or human authority; others described ways in which inputs, caveats, alternatives, dissent or lessons became less visible through summarization, ranking, omission or memory updating. We interpret these patterns as knowledge coupling, knowledge decoupling and knowledge hiding, respectively. We treat



these three as knowledge-transformation mechanisms, not governance tools or epistemic states. The epistemic conditions they create, namely, traceability, reconstructability and contestability, define what governance must protect and firm-level safeguards are responses to those conditions.

Knowledge coupling links decision-relevant knowledge to identifiable evidence, assumptions, criteria, actors and prior learning. It keeps decision rationales traceable, contestable and reusable, consistent with KM theory on integrating distributed knowledge and transforming information into actionable organizational knowledge (Grant, 1996; Nonaka, 1994; Argote and Miron-Spektor, 2011).

Knowledge decoupling weakens or breaks these links. In the accounts, it appeared both as a risk, when summaries, rankings or rationales detached from their evidentiary bases and as a governance response, when GenAI-assisted synthesis was separated from human authorization, source-of-truth validation or final justification (Bromley and Powell, 2012).

Knowledge hiding refers not only to intentional concealment by organizational actors, but also to GenAI-mediated obscuring. Retrieval, summarization, ranking or framing may foreground some information while reducing the visibility of inputs, caveats, alternatives, uncertainty, dissent or counterevidence (Connelly et al., 2012; Kellogg et al., 2020; Arias-Pérez and Vélez-Jaramillo, 2022).

However, these three patterns did not emerge with equal salience in every phase. The phase-level analysis, therefore, foregrounds the governance dynamic most strongly reflected in each DDDM activity, while noting secondary, occasional or implicit forms of knowledge coupling, knowledge decoupling and knowledge hiding where relevant.

4.2 GenAI and knowledge management across the DDDM cycle

We organize the findings into the seven DDDM phases (Cristofaro et al., 2025), which structured both the questionnaire and the analysis. The framework is an interpretive reconstruction of managerial accounts, not direct evidence that GenAI causally reorganizes governance.

Across phases, respondents rarely described GenAI as holding formal decision authority. Instead, they positioned it outside formal accountability but influential in shaping framing, evidence aggregation, option generation, ranking support, implementation guidance, performance summarization and memory updating. GenAI thus appeared quasi-agentic, shaping evidence, rationales and visibility without formal authority, intentionality or accountability. Some firms reported mature governance routines, others narrower uses for drafting, synthesis, translation or workflow acceleration.

Table 1 summarizes this phase-sensitive interpretation, aligning each DDDM phase with the dominant knowledge dynamic and related governance concern around visibility, traceability, contestability or human ownership of justification. Knowledge coupling, decoupling and hiding rarely appeared in isolation and their salience varied across firms and contexts.

4.2.1 Phase 1 – problem identification. Problem identification begins when organizational actors transform a business issue or opportunity into a decision problem by specifying objectives, constraints, success criteria, reference points and decision responsibilities. This phase matters because later evaluation, accountability and monitoring depend on how the problem is initially framed and what becomes visible as relevant evidence or legitimate concern (Jensen and Meckling, 1976; Eisenhardt, 1989).

In the managerial accounts, GenAI's role at this stage was described less in terms of evidence retrieval or option generation and more in terms of framing support. Respondents explained that GenAI helped articulate problem statements, compare alternative framings, summarize contextual information and surface preliminary assumptions. At the same time, several

Table 1 GenAI quasi-agency across DDDM phases

DDDM phase	GenAI role	Dominant knowledge dynamic	Main governance concern	Governance boundary
<i>Problem identification</i>	Helps articulate, compare and refine problem framings	Knowledge hiding through omission-by-framing	Assumptions, alternatives or concerns may disappear once a plausible frame stabilizes	Managers retain ownership of objectives, trade-offs, success criteria and final problem definition
<i>Information collection</i>	Retrieves, aggregates, summarizes and organizes evidence	Knowledge coupling through evidence aggregation	Retrieval and synthesis may hide transformations, caveats, minority signals or uncertainty	Evidence must remain traceable through traceability records, auditable prompts and versioned outputs
<i>Alternative identification</i>	Generates scenarios, counterfactuals and candidate courses of action	Knowledge coupling through generative recombination	Generated options may create premature convergence or narrow the visible option space	GenAI proposes; managers screen, justify, include, exclude and commit
<i>Alternative prioritization</i>	Compares, ranks, summarizes and structures options	Knowledge decoupling between ranking support and human justification	Rankings may become implicit organizational rationales; disagreement may disappear	Humans must author the "because" section, document dissent and retain override authority
<i>Implementation</i>	Supports operational coordination and bounded execution	Knowledge decoupling through explicit delegation boundaries	Responsibility displacement, workflow drift or hidden discretion during execution	Execution must be logged, interruptible, reviewable and constrained by nondelegable decision classes
<i>Results measurement</i>	Synthesizes performance information into dashboards, summaries and evaluative categories	Knowledge coupling through shared performance representations	Generated summaries may define salience, compress uncertainty or obscure source records	Generated narratives must remain separable from source-of-truth records and traceable metrics
<i>Results review</i>	Organizes, retrieves and reconnects prior decisions and lessons	Knowledge coupling through learning retention and memory integration	Organizational memory may become sanitized, compressed or stripped of rationale and context	Lessons must be validated, contextualized and stored with rationales, not only metrics

Source(s): Own elaboration

respondents treated this phase as particularly sensitive because early framing influenced what would later count as relevant evidence, feasible alternatives or acceptable outcomes. Some respondents also noted that GenAI's influence remained limited when managers already entered discussions with strongly established assumptions or predefined objectives.

The dominant knowledge-management issue in this phase was knowledge hiding through omission-by-framing. Unlike later phases, where traceability became more central, respondents here were primarily concerned with how assumptions, alternatives or contextual elements could quietly disappear once a plausible frame became established. One respondent captured this tension directly: "GenAI feels helpful, but it quietly changes who "owns" the frame" (118–02). Another explained that "the first framing usually shapes the whole discussion unless someone actively challenges it" (041–03).

Several respondents described governance routines intended to reduce this risk. One manager referred to a "why-not" log introduced "to reduce algorithmic knowledge hiding" (008–06). In practice, this mechanism documented excluded framings, rejected assumptions and alternatives that were considered but not pursued. Rather than treating knowledge hiding as deliberate concealment, respondents more often described it as an unintended consequence of speed, summarization or premature convergence around a persuasive framing. In this sense, the issue was less that information was intentionally hidden and more that some perspectives became less visible once a generated frame structured subsequent discussion.

Knowledge decoupling appeared more selectively in this phase as a governance response intended to preserve human ownership of the final problem definition. Several respondents described efforts to separate GenAI-assisted articulation from managerial authorization of

objectives, trade-offs and decision criteria. One respondent explained that “GenAI can help phrase the issue, but managers still need to decide what the real problem actually is” (073–04). Knowledge coupling was present only in a limited and ambivalent sense, mainly when shared framing accelerated coordination around a common interpretation of the issue.

4.2.2 Phase 2 – information collection. Information collection lays the evidentiary basis for later evaluation, prioritization and accountability. Respondents described this phase as more technically structured and operationally visible than problem identification. Here, GenAI was primarily used for retrieval, aggregation, summarization and evidence organization rather than framing support.

The dominant dynamic in this phase was knowledge coupling through evidence aggregation. Respondents frequently described GenAI as helping connect dispersed documents, policies, archives, reports and operational records into more unified evidence bases. One respondent referred to “Retrieval-Augmented Generation-based retrieval across policy and ticket archives” (023–06), while another explained that “GenAI reduced the time needed to connect information scattered across teams” (061–05). In these accounts, GenAI reorganized fragmented organizational materials into more accessible and comparable forms, improving coordination and helping managers work from a more shared evidentiary base, consistent with KM research emphasizing the integration of distributed knowledge (Kogut and Zander, 1992; Grant, 1996).

Respondents rarely treated these coupled evidence bases as self-validating. The main governance concern was that retrieval and summarization could transform evidence in ways that later became difficult to inspect. Knowledge decoupling, therefore, emerged as a governance response. One respondent described the safeguard as “auditable prompts and versioned outputs in the decision record” (049–06), while another explained that the firm had “created a traceability trail for every claim used in the decision memo” (114–06). These controls attempted to preserve visibility into how evidence had been assembled, filtered or summarized before entering managerial deliberation.

Knowledge hiding in this phase was usually described less as intentional concealment and more as hidden transformation. Respondents worried that caveats, minority signals, contextual details or uncertainty could disappear during summarization and synthesis. One manager noted that “the cleaner the synthesis becomes, the easier it is to forget what got compressed out” (028–08). Several firms therefore treated traceability as part of the evidence itself rather than as secondary documentation, consistent with AI-enabled KM research emphasizing validation and contextual verification before generated outputs become legitimate organizational knowledge (Leoni et al., 2024; Zhang et al., 2025).

Some firms supplemented technical controls with social validation mechanisms. One respondent explained that the organization had “strengthened communities of practice to validate GenAI outputs socially” (001–06), suggesting that peer review and collective interpretation remained important when assessing the credibility and contextual fit of generated outputs.

4.2.3 Phase 3 – alternatives identification. Alternatives identification concerns the construction of possible courses of action over which later evaluation, prioritization and commitment occur. Respondents described GenAI as particularly useful in this phase for expanding the option space through scenario generation, recombination of existing information, counterfactual prompting and rapid drafting of possible interventions. Compared with information collection, the focus here shifted from traceability toward the admissibility and credibility of generated alternatives.

The dominant dynamic was knowledge coupling through generative recombination. Respondents described GenAI as helping connect previously separated pieces of organizational knowledge into new candidate solutions or scenarios. One manager explained

that “GenAI was used for scenario generation and counterfactual prompting” (057–03), while another noted that “it helps us think beyond the usual alternatives we would normally discuss internally” (082–04). In these accounts, GenAI expanded the visible option space by recombining existing information into alternative strategic or operational possibilities, consistent with KM arguments that innovation depends on recombining distributed knowledge (Kogut and Zander, 1992; Grant, 1996). Some respondents added that the usefulness of generated alternatives depended on managerial expertise and prompt quality.

This expansion was not uniformly positive. Several accounts suggested that GenAI could widen perceived options while accelerating premature convergence around the first plausible set. One respondent observed that “the role inversion is strongest when speed pressures push us to accept the first output” (001–02) and another noted that “the system generates many options, but teams often stop exploring once something sounds convincing enough” (034–05). In some cases, generated alternatives became difficult to challenge because their structured presentation created an impression of completeness or analytical rigor.

Knowledge decoupling appeared mainly as a governance response separating proposal generation from managerial commitment. Generated proposals required human screening, contextual review and organizational justification before entering the formal decision set. One respondent stated explicitly, “I treat GenAI as an agent that proposes – never as an authority that decides” (114–09). We interpret this as a metaphor for proposal generation rather than evidence that GenAI becomes a contractual or intentional agent. Some managers also introduced manually generated alternatives to avoid excessive dependence on AI-supported option sets.

Knowledge hiding appeared less systematically than in information collection. Several respondents worried that alternatives not surfaced by GenAI could disappear from deliberation, narrowing the option space rather than concealing evidence. Some firms responded with documented override rights and rationales for exclusions.

4.2.4 Phase 4 – alternatives prioritization. Alternatives prioritization concerns the evaluation and ordering of admissible options against organizational criteria, trade-offs, risks and uncertainty. Respondents described GenAI here less as a generator of alternatives and more as a comparative synthesis device that helped sort, rank, summarize and structure competing options. Compared with earlier phases, managerial concerns shifted toward justificatory ownership, namely, whether GenAI-supported rankings might become the implicit rationale for organizational commitment.

The dominant dynamic was knowledge decoupling between system-assisted prioritization and human-authored justification. Respondents described GenAI as useful for reducing informational complexity and making competing alternatives more comparable. One manager explained that GenAI helped with “triaging noisy signals into ranked themes” (103–09). Ranking support accelerated comparison and reduced cognitive overload during evaluation, particularly under high information volume, compressed timelines or limited managerial attention. Respondents emphasized that comparative ordering alone could not constitute a sufficient organizational rationale for action.

Several firms introduced governance routines to preserve human ownership of the final justification. One respondent described the safeguard directly, stating that “we prevent delegated authorship by requiring humans to write the “because” section” (018–01). The “because” section documented why an option was ultimately selected and who remained accountable for the decision rationale. Another manager similarly explained that “GenAI can rank priorities, but managers still need to explain why this option deserves organizational backing” (071–02).

Respondents also described prioritization support as simultaneously useful and fragile. Ranking systems reduced informational complexity but could create a false sense of

stability or consensus around one ordering. One respondent captured this tension by observing that “the agent role flips when GenAI’s narrative becomes the default justification” (006–10). We interpret this wording as a respondent metaphor for the displacement of justificatory authorship rather than as evidence that GenAI acquires intentionality, accountability or formal authority. Quantified rankings often appeared more authoritative internally, even when managers remained uncertain about the underlying assumptions.

Knowledge hiding emerged mainly through suppressed uncertainty, reduced visibility of disagreement and overconfidence in ranked outputs. One manager explained, “we rely on thresholds for escalation when uncertainty is high. And that’s why we document dissent, not just outcomes” (005–04). Some organizations preserved contestability by documenting disagreement and unresolved concerns rather than recording only the final ranking, though dissent documentation was uneven across teams. Knowledge coupling remained secondary, beneficial only when prioritization stayed open to challenge.

4.2.5 Phase 5 – alternative implementation. Alternative implementation concerns translating a selected course of action into operational execution through workflows, task allocation, escalation paths and coordination routines. In classic agency theory, implementation is particularly sensitive because organizational actors exercise discretion under conditions in which effort, interpretation and local adaptation are only partially observable (Jensen and Meckling, 1976; Eisenhardt, 1989). Respondents generally did not describe this phase as full automation. Instead, they described forms of bounded delegation in which GenAI could support implementation only within constrained, logged, reviewable and interruptible operational environments.

The dominant dynamic in this phase was knowledge decoupling through explicit delegation boundaries. Respondents emphasized that GenAI-supported execution remained acceptable only when organizations maintained visible limits around what systems could operationally trigger, modify, authorize or commit. One manager summarized this governance logic as “rules, logs, overrides” (049–08). Across the accounts, GenAI became organizationally consequential not through autonomous authority, but through sociotechnical arrangements specifying where intervention was permitted, reviewable and reversible.

Several respondents described governance mechanisms intended to preserve these boundaries. The clearest examples involved nondelegable decision classes, escalation requirements, override rights and human sign-off procedures. One respondent referred to “explicit nondelegable decision classes” (114–02), explaining that decisions involving legal exposure, reputational implications or strategic commitment could not be operationally executed without managerial authorization. Another manager noted that “automation is acceptable only when someone can still interrupt the process without friction” (038–03). These accounts suggest that implementation governance focused less on whether GenAI could execute tasks and more on preserving visible accountability boundaries around execution authority.

Some respondents described formalized controls involving escalation paths, execution logs, approval chains and audit visibility, while others described looser arrangements where GenAI remained closer to workflow assistance than delegated action. One respondent explained that the system was mainly used for “meeting-note synthesis and action-item extraction,” adding that “we keep the system in the loop, but we keep responsibility on the human side” (003–05). Another described GenAI as “helpful for operational coordination, but not trusted enough for autonomous execution” (064–02). These accounts complicate the bounded-delegation pattern by showing substantial variation in how deeply firms operationally embedded GenAI into execution processes.

Knowledge coupling remained present but secondary. Respondents described implementation knowledge becoming embedded in prompts, workflow configurations,

escalation protocols, reporting templates and operational rules coordinating recurring activities across teams. Knowledge hiding appeared more intermittently and was usually associated with responsibility displacement or silent workflow drift rather than with concealment of evidence. Several respondents worried that managers could gradually lose visibility into where discretionary intervention still occurred once execution processes became highly system-supported. One respondent described the safeguard as “a human signature requirement on final rationales” (114–01), explaining that “GenAI can compute, but it cannot legitimately commit on our behalf.”

4.2.6 Phase 6 – results measurement. Results measurement concerns the evaluation of realized outcomes against objectives and the rationale that justified commitment, determining how performance becomes visible and comparable. Respondents described GenAI-supported measurement as a process that shaped how actors interpreted performance signals, categorized outcomes and defined success, deviation or failure.

The dominant dynamic was knowledge coupling through shared performance representations and evaluative categories. Respondents described GenAI as useful for synthesizing dispersed operational information into coherent dashboards, summaries and comparative views. One manager explained that “GenAI acts like a principal by setting the categories and the salience map” (034–09). We interpret this wording as a respondent metaphor for epistemic influence rather than evidence that GenAI formally becomes a principal. The account illustrates a recurring concern that GenAI-supported measurement could influence which indicators became salient and which outcomes attracted attention. Another similarly noted that “once the dashboard stabilizes, people start treating the categories as reality itself” (021–07). Managers became increasingly dependent on dashboards in complex or data-intensive environments.

Knowledge decoupling appeared through attempts to separate generated summaries from validated underlying records. Respondents emphasized that synthesized narratives could not substitute for the underlying data that produced them. One respondent noted that “we added “source-of-truth” tags to prevent synthetic text from becoming canon” (032–02). Another explained that “managers must still be able to trace summaries back to the originating metrics and records” (087–05). These practices preserved inspectability by distinguishing generated interpretations from validated evidence. Verification routines were difficult to sustain when reporting cycles accelerated or information volume increased.

Knowledge hiding emerged as the loss of uncertainty, contextual nuance or minority signals during summarization. Respondents worried that fluent narratives could compress caveats and conflicting indicators into cleaner but less contestable representations. One manager referred to “model choice rules that privilege transparency over marginal accuracy” (114–09), explaining that the firm preferred less predictive sophistication, allowing managers to inspect assumptions and classifications. Visually polished summaries sometimes generated stronger internal confidence than the underlying data quality justified.

Governance approaches varied considerably. Some firms described formalized controls involving traceability, model governance, dashboard validation and source verification, while others relied on managerial judgment. One respondent explained that “we still cross-check important summaries manually because we do not fully trust narrative compression” (055–03). Another described the system as “helpful for seeing patterns quickly, but risky if people stop questioning how the patterns were constructed” (012–08).

4.2.7 Phase 7 – results review. Results review concerns how organizations interpret realized outcomes, compare them against earlier expectations and rationales and translate experience into future decision processes. Respondents described this phase as more than retrospective evaluation or performance reporting. Instead, they treated it as the stage at which organizational experience was retained, reformulated and embedded in future routines, governance arrangements and decision-making practices. In this sense, the

review connected one DDDM cycle to the next by influencing what future decision-makers could later retrieve, reuse, challenge or treat as organizationally legitimate knowledge, consistent with organizational learning research emphasizing retention and incorporation into organizational memory systems (Argote and Miron-Spektor, 2011).

The dominant dynamic in this phase was knowledge coupling through learning retention and memory integration. Several respondents described GenAI as useful for organizing, retrieving and reconnecting prior decision experiences across projects, teams or time periods. One respondent explained that “GenAI was used for semantic search over lessons-learned repositories” (017–09), while another noted that “the system helps reconnect current decisions with older cases we would probably not remember otherwise” (026–04). In these accounts, GenAI-supported review increased the portability and retrievability of prior organizational experience by reconnecting historical rationales, contextual lessons and implementation experiences to future decision situations.

At the same time, respondents did not describe organizational memory as a neutral or automatic repository. Knowledge decoupling appeared through governance routines separating local experience from formal organizational-memory updating. One manager explained that “we redesigned how lessons move from local teams to corporate memory” (009–07), suggesting that not all experiential knowledge was automatically institutionalized. Several respondents described review as a translation process in which lessons had to be validated, contextualized and made organizationally reusable before entering more stable repositories or governance systems.

Knowledge hiding appeared mainly through the risk of sanitized or overly compressed memory. Respondents worried that failures, contextual caveats, uncertainty or dissenting interpretations could gradually disappear when review outputs were reduced to dashboards, summaries or polished retrospective narratives. One respondent explained that “we updated our knowledge base after-action, not just our dashboards” (001–01), while another noted that the organization had to “store rationales, not just metrics, to preserve decision lineage” (005–03). These accounts suggest that preserving interpretive context mattered as much as storing outcomes themselves.

Some respondents described structured review practices involving after-action documentation, searchable repositories, formal lessons-learned systems and retrospective governance reviews. Others described looser arrangements in which organizational learning remained informal, fragmented or dependent on individual managerial initiative. One respondent observed that “a lesson only survives if someone actively reuses it later” (072–02), highlighting that retention alone did not guarantee organizational learning.

5. Discussion


The phase-level findings show that GenAI’s role varies substantially across the DDDM cycle. Its influence was limited in peripheral administrative tasks and stronger when outputs shaped problem framing, evidence construction, prioritization, implementation control, performance interpretation or organizational memory. This pattern matters for principal–agent theory because GenAI becomes consequential precisely where data are transformed into evidence, rationales, classifications and memory structures. Agency problems therefore extend beyond hidden action, effort observability and incentive misalignment toward hidden knowledge transformation, namely, how evidence is selected, summarized, compressed, ranked, reframed and stabilized as legitimate grounds for action (Jensen and Meckling, 1976; Eisenhardt, 1989; Kellogg *et al.*, 2020; Cristofaro and Bañón-Gomis, 2026). More fundamentally, GenAI begins to reconfigure what counts as legitimate organizational knowledge, since legitimacy increasingly depends on machine-mediated transformations that no single human fully authors or reconstructs.

Building on research on digital and artificial agency (Ågerfalk, 2020; Dattathrani and De, 2023) and AI-related control and accountability (Humberd and Latham, 2025), we define GenAI quasi-agency as a form of epistemic mediation with agency-relevant consequences. GenAI becomes quasi-agentive when it materially shapes framing, evidence construction, prioritization, justification, performance representations or organizational memory while lacking formal authority, intentionality, autonomous goals and accountability. The construct does not imply that GenAI becomes a contractual, moral or intentional actor. It captures the narrower condition in which GenAI remains formally subordinate to human decision-makers while influencing the epistemic conditions under which authority is exercised.

Classic agency theory treats information asymmetry primarily as hidden action, whereas in GenAI-supported settings asymmetry increasingly concerns hidden knowledge transformation (Jensen and Meckling, 1976; Eisenhardt, 1989). Governance therefore extends from monitoring behavior and outcomes toward governing how evidence becomes rationale organizational memory and legitimate grounds for action under epistemic opacity. We adopt the quasi prefix to mark a partial and incomplete form of agency in which GenAI exercises agency-relevant epistemic influence over framing, evidence and justification while lacking intentionality, autonomous goals, formal authority and accountability. We prefer it to digital agency, which foregrounds enactment through code and sociotechnical arrangements and to artificial agency, which foregrounds autonomy and adaptivity, because our concern is the as-if agency that emerges when generated outputs shape the epistemic conditions of decision-making without the system owning or answering for the decision. This is also what separates quasi-agency from a relabeling of epistemic influence, since GenAI shapes the conditions under which authority is exercised, so agency-relevant governance concerns arise before any formal decision is taken rather than only at the moment of choice (Ågerfalk, 2020; Dattathrani and De, 2023).

Figure 2 summarizes three uneven knowledge dynamics across the DDDM cycle (knowledge coupling, knowledge decoupling and knowledge hiding). Knowledge coupling stabilizes shared interpretations by integrating fragmented evidence, assumptions, rationales and prior learning into portable decision representations. Knowledge decoupling has a dual role. As a risk, outputs may detach from traceability, assumptions, uncertainty or contextual grounding. As a governance response, organizations preserve human justificatory ownership through

Figure 2 The GenAI quasi-agency knowledge governance framework



traceability controls, override routines, escalation mechanisms and separation between machine-assisted synthesis and human authorization. Knowledge hiding includes both deliberate concealment and GenAI-mediated obscuring through omission, filtering, ranking, compression or loss of caveats. On the governance side, the figure identifies the epistemic conditions these dynamics place at stake, namely, traceability, reconstructability and contestability and frames human justificatory ownership as the accountability these conditions exist to protect. Together, these dynamics shape what becomes visible, credible, reconstructable, contestable and actionable in organizational decision-making (Connelly *et al.*, 2012; Kellogg *et al.*, 2020; Leoni *et al.*, 2024).

These mechanisms did not appear uniformly across the DDDM cycle. Framing sovereignty dominated problem identification, traceability and evidence traceability dominated information collection, justificatory ownership became most salient during prioritization, bounded delegation characterized implementation and learning retention and organizational memory became central during review. The framework should therefore be interpreted as an abductive reconstruction of uneven governance dynamics rather than a deterministic causal model, consistent with the study's interpretive design (Braun and Clarke, 2019). The cycle is therefore a theoretical requirement rather than an organizing convenience, because agency-relevant risk changes with the epistemic operation each phase performs, so only a phase model shows that governance requirements are heterogeneous rather than constant.

GenAI should not be treated as quasi-agentic merely because it appears in managerial workflows. Its relevance depends on whether it shapes decision-relevant knowledge transformations rather than supporting peripheral tasks. When used for formatting, translation, grammar correction or low-risk administrative support, it remains closer to a productivity tool. Its quasi-agentic significance increases when outputs shape framing, evidence selection, synthesis, ranking, justification, performance categorization or organizational memory, particularly when these outputs circulate as decision-ready representations under limited reconstructability and weak human challenge routines. Deployment mode also matters because general-purpose chatbots, enterprise copilots, retrieval-augmented systems and proprietary applications differ substantially in traceability, controllability, explainability and auditability (Leoni *et al.*, 2024; Hillebrand *et al.*, 2025).

Alternative explanations remain plausible. Some observed routines may reflect broader digital maturity, workflow automation, documentation practices, regulatory pressure or managerial fashion rather than GenAI alone (Abrahamson, 1996). Our claim is therefore deliberately circumscribed. GenAI becomes theoretically consequential when it materially participates in transforming evidence into decision-relevant knowledge, justification and organizational memory.

6. Implications and conclusions

This study makes three theoretical contributions. First, it extends KM research by shifting attention from information access toward epistemic legitimacy in GenAI-supported decision-making. KM research has long distinguished information artifacts from actionable knowledge, valuable only when interpreted, contextualized and legitimized in practice (Nonaka, 1994; Grant, 1996). GenAI intensifies the fragility of this distinction because fluent, decision-ready outputs can acquire unwarranted legitimacy even when traceability, assumptions and uncertainty remain difficult to reconstruct (Zhang *et al.*, 2025). Three interconnected mechanisms develop this contribution. Knowledge coupling integrates fragmented information into shared summaries and narratives that coordinate distributed actors, knowledge decoupling captures both the risk that outputs detach from their evidentiary basis and the safeguards that separate synthesis from human accountability and knowledge hiding extends beyond deliberate concealment to the disappearance of caveats or dissent during synthesis (Connelly *et al.*, 2012; Arias-Pérez and Vélez-Jaramillo,

2022). Together, these mechanisms reposition KM under GenAI as an epistemic governance problem centered on preserving traceability, reconstructability, contestability and human accountability.

Second, the study extends agency theory beyond hidden action, information asymmetry and incentive alignment (Jensen and Meckling, 1976; Eisenhardt, 1989). Organizational actors increasingly receive coherent rankings and rationales without reconstructing how evidence was selected, reframed or legitimized. Agency problems therefore concern not only the delegation of action but also the delegation of knowledge transformation. Building on digital and artificial agency research (Ågerfalk, 2020; Dattathrani and De, 2023; Humberd and Latham, 2025), GenAI quasi-agency names the capacity to shape framing, evidence, justification and memory without possessing formal authority, intentionality or accountability. GenAI does not independently commit the organization to action. Its influence operates through epistemic mediation, affecting how organizational actors interpret information, construct rationales and legitimize decisions.

Third, the study contributes to DDDM research by showing that GenAI's influence is uneven across decision activities, so that GenAI should be treated not as a generic decision-support technology but as a phase-sensitive epistemic infrastructure whose governance requirements vary across the decision cycle (Brynjolfsson *et al.*, 2011).

From a practical perspective, organizations should calibrate governance intensity according to GenAI's epistemic role rather than adoption alone. Low-risk uses require minimal oversight, while stronger safeguards are necessary when GenAI contributes to evidence selection, prioritization, justification, performance categorization or organizational memory. Concrete mechanisms, set out by phase in Table 1, range from “why-not” logs and traceability tags to human-authored “because” sections, nondelegable decision classes and after-action reviews that retain rationales rather than only outcomes. Many of these safeguards remain feasible even for SMEs without sophisticated AI-governance infrastructures, because effective GenAI governance depends less on technological sophistication than on preserving visibility, contestability and human accountability across knowledge transformations.

The study has limitations that bound its claims. The empirical setting focuses on Italian SMEs already using GenAI in recurring workflows, which may transfer less directly to large firms, highly regulated sectors or public organizations. The study relies on structured written responses from one respondent per firm, limiting access to interactional dynamics, disagreement and political processes. The analysis does not include behavioral traces, prompt histories or retrieval logs and therefore cannot verify how outputs were generated or whether governance routines operated as described. The cross-sectional design cannot capture the evolution of governance routines over time. In addition, self-reported accounts of governance routines may be subject to social desirability bias, as respondents may describe more mature or more consistently enacted oversight practices than those actually followed in their firms. This consideration reinforces the interpretation of the framework as an abductive reconstruction of reported practice rather than as independently verified behavior.

Future research should examine variation across deployment modes, which differ in traceability, controllability and auditability. Comparative work across firm sizes, industries and levels of digital maturity could identify when traceability, contestability and human justificatory ownership become harder to preserve. Longitudinal studies could clarify whether GenAI governance stabilizes into formal infrastructures or remains dependent on local improvisation. Additional research should investigate the cognitive and social consequences of prolonged interaction with GenAI outputs, including how trust attribution and epistemic dependence shape reliance on generated rationales, and how organizations preserve tacit and experiential knowledge when rationales are synthesized through generative systems.

GenAI is becoming embedded in how organizations construct evidence, justify decisions, evaluate outcomes and retain organizational memory. Its relevance extends beyond automation or productivity enhancement. By introducing GenAI quasi-agency and developing a phase-sensitive governance framework, this study shows that organizational risks emerge not only from hidden human action but also from hidden knowledge transformation. The future of AI governance will depend less on replacing human judgment and more on governing the epistemic infrastructures through which organizational knowledge is transformed, stabilized and legitimized.

References

- Abrahamson, E. (1996), "Management fashion", *Academy of Management Review*, Vol. 21 No. 1, pp. 254-285.
- Ågerfalk, P.J. (2020), "Artificial intelligence as digital agency", *European Journal of Information Systems*, Vol. 29 No. 1, pp. 1-8, doi: [10.1080/0960085x.2020.1721947](https://doi.org/10.1080/0960085x.2020.1721947).
- Arias-Pérez, J. and Vélez-Jaramillo, J. (2022), "Understanding knowledge hiding under technological turbulence caused by artificial intelligence and robotics", *Journal of Knowledge Management*, Vol. 26 No. 6, pp. 1476-1491, doi: [10.1108/JKM-01-2021-0058](https://doi.org/10.1108/JKM-01-2021-0058).
- Argote, L. and Miron-Spektor, E. (2011), "Organizational learning: from experience to knowledge", *Organization Science*, Vol. 22 No. 5, pp. 1123-1137, doi: [10.1287/orsc.1100.0621](https://doi.org/10.1287/orsc.1100.0621).
- Arnaut, B., El Nemer, S., Kokkinopoulou, E., Skaf, Y., Vrontis, D., Esposito, M. and Rebeiz, K.S. (2026), "Bridging AI and human intelligence: advancing knowledge ecosystems through effective communication for strategic innovation in the US utility industry", *Journal of Knowledge Management*, Vol. 30 No. 4, pp. 1468-1488, doi: [10.1108/JKM-05-2025-0681](https://doi.org/10.1108/JKM-05-2025-0681).
- Bendickson, J., Muldoon, J., Liguori, E.W. and Davis, P.E. (2016), "Agency theory: background and epistemology", *Journal of Management History*, Vol. 22 No. 4, pp. 437-449, doi: [10.1108/JMH-06-2016-0028](https://doi.org/10.1108/JMH-06-2016-0028).
- Braun, V. and Clarke, V. (2006), "Using thematic analysis in psychology", *Qualitative Research in Psychology*, Vol. 3 No. 2, pp. 77-101, doi: [10.1191/1478088706qp0630a](https://doi.org/10.1191/1478088706qp0630a).
- Braun, V. and Clarke, V. (2019), "Reflecting on reflexive thematic analysis", *Qualitative Research in Sport, Exercise and Health*, Vol. 11 No. 4, pp. 589-597, doi: [10.1080/2159676X.2019.1628806](https://doi.org/10.1080/2159676X.2019.1628806).
- Bromley, P. and Powell, W.W. (2012), "From smoke and mirrors to walking the talk: decoupling in the contemporary world", *Academy of Management Annals*, Vol. 6 No. 1, pp. 483-530, doi: [10.5465/19416520.2012.684462](https://doi.org/10.5465/19416520.2012.684462).
- Brynjolfsson, E., Hitt, L.M. and Kim, H.H. (2011), "Strength in numbers: how does data-driven decision-making affect firm performance?", SSRN Working Paper, p. 1819486, doi: [10.2139/ssrn.1819486](https://doi.org/10.2139/ssrn.1819486).
- Chen, Q.Q., Lin, L.M. and Liu, M. (2026), "Enhancing knowledge sharing in generative AI integration: the impact of AI self-efficacy and skill threat perceptions", *Journal of Knowledge Management*, Vol. 30 No. 3, pp. 1077-1100, doi: [10.1108/JKM-11-2024-1328](https://doi.org/10.1108/JKM-11-2024-1328).
- Connelly, C.E., Zweig, D., Webster, J. and Trougakos, J.P. (2012), "Knowledge hiding in organizations", *Journal of Organizational Behavior*, Vol. 33 No. 1, pp. 64-88, doi: [10.1002/job.737](https://doi.org/10.1002/job.737).
- Cristofaro, M. and Bañón-Gomis, A.J. (2026), "Dancing with the algorithm: a framework to navigate knowledge and autonomy in AI-assisted managerial decisions", *Journal of Knowledge Management*, Vol. 30 No. 11, pp. 1-31, doi: [10.1108/JKM-06-2025-0870](https://doi.org/10.1108/JKM-06-2025-0870).
- Cristofaro, M. and Giardino, P.L. (2025), "Surfing the AI waves: the historical evolution of artificial intelligence in management and organizational studies and practices", *Journal of Management History*, Vol. 32 No. 4, pp. 670-696, doi: [10.1108/JMH-01-2025-0002](https://doi.org/10.1108/JMH-01-2025-0002).
- Cristofaro, M., Giardino, P.L. and Barboni, L. (2025), "Growth hacking: a scientific approach for data-driven decision making", *Journal of Business Research*, Vol. 186, p. 115030, doi: [10.1016/j.jbusres.2024.115030](https://doi.org/10.1016/j.jbusres.2024.115030).
- Dattatharani, S. and De, R. (2023), "The concept of agency in the era of artificial intelligence: dimensions and degrees", *Information Systems Frontiers*, Vol. 25 No. 1, pp. 29-54, doi: [10.1007/s10796-022-10336-8](https://doi.org/10.1007/s10796-022-10336-8).

- Eisenhardt, K.M. (1989), "Agency theory: an assessment and review", *The Academy of Management Review*, Vol. 14 No. 1, pp. 57-74, doi: [10.2307/258191](https://doi.org/10.2307/258191).
- Gioia, D.A., Corley, K.G. and Hamilton, A.L. (2013), "Seeking qualitative rigor in inductive research: notes on the Gioia methodology", *Organizational Research Methods*, Vol. 16 No. 1, pp. 15-31, doi: [10.1177/1094428112452151](https://doi.org/10.1177/1094428112452151).
- Glaser, B. and Strauss, A. (1967), *The Discovery of Grounded Theory: Strategies for Qualitative Research*, Aldine.
- Grant, R.M. (1996), "Toward a knowledge-based theory of the firm", *Strategic Management Journal*, Vol. 17 No. S2, pp. 109-122, doi: [10.1002/smj.4250171110](https://doi.org/10.1002/smj.4250171110).
- Hernández-Tamurejo, Á., González-Padilla, P. and Saiz-Sepúlveda, Á. (2025), "The economics of AI adoption in OTAs: market dynamics and future research", *Global Economics Research*, Vol. 1 No. 1, p. 100001, doi: [10.1016/j.ecores.2025.100001](https://doi.org/10.1016/j.ecores.2025.100001).
- Hillebrand, L., Raisch, S. and Schad, J. (2025), "Managing with artificial intelligence: an integrative framework", *Academy of Management Annals*, Vol. 19 No. 1, pp. 343-375, doi: [10.5465/annals.2022.0072](https://doi.org/10.5465/annals.2022.0072).
- Humberd, B.K. and Latham, S.F. (2025), "When AI becomes an agent of the firm: examining the evolution of AI in organizations through an agency theory lens", *Journal of Management Studies*, Vol. 63 No. 2, doi: [10.1111/joms.13274](https://doi.org/10.1111/joms.13274).
- Jensen, M.C. and Meckling, W.H. (1976), "Theory of the firm: managerial behavior, agency costs and ownership structure", *Journal of Financial Economics*, Vol. 3 No. 4, pp. 305-360, doi: [10.1016/0304-405X\(76\)90026-X](https://doi.org/10.1016/0304-405X(76)90026-X).
- Kampoowale, I. (2025), "Linking big data analytics capabilities to organizational learning through knowledge management and data-driven decision-making", *The TQM Journal*, Vol. 38 No. 6, pp. 1328-1347, doi: [10.1108/TQM-08-2024-0282](https://doi.org/10.1108/TQM-08-2024-0282).
- Kellogg, K.C., Valentine, M.A. and Christin, A. (2020), "Algorithms at work: the new contested terrain of control", *Academy of Management Annals*, Vol. 14 No. 1, pp. 366-410, doi: [10.5465/annals.2018.0174](https://doi.org/10.5465/annals.2018.0174).
- Kogut, B. and Zander, U. (1992), "Knowledge of the firm, combinative capabilities, and the replication of technology", *Organization Science*, Vol. 3 No. 3, pp. 383-397, doi: [10.1287/orsc.3.3.383](https://doi.org/10.1287/orsc.3.3.383).
- Langley, A. (1999), "Strategies for theorizing from process data", *The Academy of Management Review*, Vol. 24 No. 4, pp. 691-710, doi: [10.5465/amr.1999.2553248](https://doi.org/10.5465/amr.1999.2553248).
- Leoni, L., Gueli, G., Ardolino, M., Panizzon, M. and Gupta, S. (2024), "AI-empowered KM processes for decision-making: empirical evidence from worldwide organisations", *Journal of Knowledge Management*, Vol. 28 No. 11, pp. 320-347, doi: [10.1108/JKM-03-2024-0262](https://doi.org/10.1108/JKM-03-2024-0262).
- Li, B., Wang, L., Fu, S. and Wan, J. (2025), "How does artificial intelligence empower enterprise technological innovation: from the perspective of knowledge management capability", *Journal of Knowledge Management*, pp. 1-31, doi: [10.1108/JKM-03-2025-0345](https://doi.org/10.1108/JKM-03-2025-0345).
- Li, J. and Yuan, Q. (2026), "The impact of AI-generated knowledge on user knowledge avoidance in OKCs: the moderating role of AI literacy and platform trust", *Journal of Knowledge Management*, Vol. 30 No. 2, pp. 1-25, doi: [10.1108/JKM-02-2025-0231](https://doi.org/10.1108/JKM-02-2025-0231).
- Lincoln, Y.S. and Guba, E.G. (1985), *Naturalistic Inquiry*, Sage, Beverly Hills, Calif.
- Mårtensson, M. (2000), "A critical review of knowledge management as a management tool", *Journal of Knowledge Management*, Vol. 4 No. 3, pp. 204-216, doi: [10.1108/13673270010350002](https://doi.org/10.1108/13673270010350002).
- Miles, M.B., Huberman, A.M. and Saldana, J. (2014), *Qualitative Data Analysis*, Sage.
- Nonaka, I. (1994), "A dynamic theory of organizational knowledge creation", *Organization Science*, Vol. 5 No. 1, pp. 14-37, doi: [10.1287/orsc.5.1.14](https://doi.org/10.1287/orsc.5.1.14).
- Stelmaszak, M., Joshi, M. and Constantiou, I. (2026), "Artificial intelligence as an organizing capability arising from human-algorithm relations", *Journal of Management Studies*, Vol. 63 No. 2, pp. 335-365, doi: [10.1111/joms.70003](https://doi.org/10.1111/joms.70003).
- Trio, O., Caboni, F., Cavallo, F. and Meissner, D. (2026), "Artificial intelligence and tacit knowledge in radio content production", *Journal of Knowledge Management*, Vol. 30 No. 4, pp. 1282-1302, doi: [10.1108/JKM-06-2025-0881](https://doi.org/10.1108/JKM-06-2025-0881).
- Vanneste, B.S. and Puranam, P. (2025), "Artificial intelligence, trust, and perceptions of agency", *Academy of Management Review*, Vol. 50 No. 4, doi: [10.5465/amr.2022.0041](https://doi.org/10.5465/amr.2022.0041).

Zhang, Q., Zuo, J. and Yang, S. (2025), "Research on the impact of generative artificial intelligence (GenAI) on enterprise innovation performance: a knowledge management perspective", *Journal of Knowledge Management*, Vol. 29 No. 7, pp. 2238-2257, doi: [10.1108/JKM-10-2024-1198](https://doi.org/10.1108/JKM-10-2024-1198).

Further reading

Yin, J. and Hoang, K.D. (2026), "AI-enabled knowledge renewal: the role of leaders' AI attitudes and unlearning in enhancing employees' creative performance", *Journal of Knowledge Management*, Vol. 30 No. 1, pp. 211-230, doi: [10.1108/JKM-02-2025-0209](https://doi.org/10.1108/JKM-02-2025-0209).

Supplementary material

The supplementary material for this article can be found online.

Corresponding author

Matteo Cristofaro can be contacted at: matteo.cristofaro@uniroma2.it

For instructions on how to order reprints of this article, please visit our website:
www.emeraldgrouppublishing.com/licensing/reprints.htm
Or contact us for further details: permissions@emeraldinsight.com