Roman KUZIV[1],
postgraduate student at the Department of Management,
ORCID: https://orcid.org/0009-0007-7984-6162

[1] *Kyiv National Economic University named after Vadym Hetman*

Прийняття: 09/03/2026
Рецензія: 17/03/2026
Публікація: 31/03/2026

JEL Класифікатор:
O33, M15, L86

ISSN 2786-5339 (print)
ISSN 2786-5347 (online)

DOI: https://doi.org/10.53920/ES-2026-1-18

BARRIERS TO AGENTIC AI ENTERPRISE TRANSFORMATION

This study identifies, classifies, and critically analyzes barriers to enterprise transformation under the influence of agentic AI – autonomous software that leverages large language models (LLMs) to perceive its environment, reason through complex tasks, plan and execute actions, and use tools to achieve goals with minimal human oversight.

A critical narrative literature review of 30 sources (2019–2026) was conducted. Barriers were identified inductively through open and axial coding; Sociotechnical Systems (STS) theory was then applied as an interpretive lens to map the resulting dimensions onto social and technical subsystems and analyze cross-subsystem interactions.

Twenty-nine barriers were classified into five dimensions: technological (7), organizational (7), human (6), governance and regulatory (4), and economic (5). Each barrier was assessed for agentic specificity. Three barriers were identified as agentic-specific (error propagation in multi-agent systems, role ambiguity, accountability diffusion). At the same time, the remaining 26 are carried over from prior digital transformation waves – 22 in amplified form and 4 unchanged. STS mapping based on root-cause analysis revealed that 12 barriers originate in the technical subsystem and 17 in the social subsystem, with governance serving as the social subsystem's primary mechanism for managing the technical subsystem. Five interaction mechanisms were identified, with the majority propagating across the subsystem boundary.

Agentic AI transformation barriers constitute an interdependent sociotechnical system rather than isolated obstacles. The governance calibration problem – balancing

control with the autonomy that gives agentic AI its value — emerges as the STS joint optimization challenge: governance, as the social subsystem's mechanism for managing the technical subsystem, must simultaneously enable and constrain autonomous operation.

The taxonomy provides a diagnostic framework for identifying priority barrier dimensions and understanding cross-dimensional amplification mechanisms. The agentic-specificity classification helps organizations distinguish challenges that require novel approaches from those that are addressable with established practices.

Keywords: *agentic AI, enterprise transformation, barriers, sociotechnical systems, AI adoption, governance, digital transformation, organizational change, taxonomy, multi-agent systems.*

Роман КУЗІВ

БАР'ЄРИ АГЕНТНОЇ ШІ-ТРАНСФОРМАЦІЇ ПІДПРИЄМСТВ

Дослідження спрямоване на ідентифікацію, класифікацію та критичний аналіз бар'єрів трансформації підприємств під впливом агентного штучного інтелекту — автономного програмного забезпечення, що використовує великі мовні моделі (LLM) щодо сприйняття середовища, аналізу складних завдань, планування та виконання дій і використання інструментів для досягнення цілей з мінімальним людським контролем.

Проведено критичний наративний огляд літератури на основі 30 джерел (2019 – 2026 рр.). Бар'єри ідентифіковано індуктивно шляхом відкритого та осьового кодування; теорію соціотехнічних систем (СТС) застосовано як інтерпретаційну лінзу для відображення виявлених вимірів на соціальну та технічну підсистеми та аналізу міжсистемних взаємодій.

Класифіковано 29 бар'єрів за п'ятьма вимірами: технологічний (7), організаційний (7), людський (6), управлінсько-регуляторний (4) та економіко-фінансовий (5). Кожен бар'єр оцінено за агентною специфічністю. Три бар'єри визначено як агентно-специфічні (каскадне поширення помилок у мультиагентних системах, невизначеність ролей, дифузія відповідальності), тоді як решта 26 є успадкованими з попередніх хвиль цифрової трансформації — 22 у підсиленій формі та 4 без змін. СТС-відображення на основі аналізу першопричин виявило, що 12 бар'єрів походять з технічної підсистеми та 17— із соціальної, при цьому управління виступає механізмом соціальної підсистеми для регулювання технічної. Ідентифіковано п'ять механізмів взаємодії, більшість з яких перетинає межу між підсистемами.

Бар'єри трансформації під впливом агентного ШІ складають взаємозалежну соціотехнічну систему, а не набір ізольованих перешкод. Проблема калібрування управління — балансування контролю з автономністю — визначена як проблема спільної оптимізації СТС, що виникає з аналізу міжсистемних взаємодій: управління, як механізм соціальної підсистеми для регулювання технічної, має одночасно уможливлювати та обмежувати автономне функціонування.

Таксономія надає діагностичний інструмент для визначення пріоритетних напрямів втручання та розуміння механізмів міжвимірного підсилення бар'єрів. Класифікація за агентною специфічністю допомагає організаціям відрізняти виклики, що потребують нових підходів, від тих, що вирішуються усталеними практиками.

Ключові слова: *агентний штучний інтелект, трансформація підприємств, бар'єри, соціотехнічні системи, впровадження ШІ, управління, цифрова трансформація, організаційні зміни, таксономія, мультиагентні системи.*

Problem statement. Enterprise adoption of artificial intelligence has accelerated from experimental pilots to a strategic priority within a single business cycle. According to industry surveys, 30% of organizations are exploring agentic AI options and 38% are piloting solutions. Yet, only 11% have deployed agentic systems in production environments (Deloitte, 2026)—this disparity between exploration and deployment points to systematic barriers that existing frameworks have not adequately addressed.

Agentic AI represents a qualitative departure from prior AI implementations. Unlike conventional AI systems that execute predefined tasks, agentic AI consists of autonomous multi-agent systems capable of dynamic task decomposition, inter-agent coordination, persistent memory, and continuous scope expansion (Sapkota et al., 2025). These properties create transformation challenges that do not map cleanly onto established technology adoption frameworks designed for bounded, predictable systems.

The academic treatment of barriers to agentic AI enterprise transformation remains fragmented. Existing analyses either treat agentic AI as generic AI, applying established frameworks such as the Technology-Organization-Environment (TOE) model (Tornatzky & Fleischer, 1990) without examining what distinguishes agentic contexts, or produce ad-hoc barrier lists without theoretical grounding (Chopra, 2025; Deloitte, 2026). Meanwhile, academic studies that do apply theoretical rigor — such as Papagiannidis et al.'s (2022) multi-case analysis of AI governance or Makarius et al.'s (2020) sociotechnical framework for AI integration – address AI broadly rather than the specific challenges of autonomous, scope-expanding agent systems.

This gap has practical consequences. Organizations attempting agentic AI transformation lack a systematic diagnostic tool for identifying which barriers require attention and how barriers in different dimensions interact. Without such a tool, interventions target isolated symptoms – investing in infrastructure while governance gaps persist, or training employees while organizational structures remain unchanged – producing the persistent pilot-to-scale failures documented by industry data.

This study addresses the gap by constructing a five-dimensional barrier taxonomy through inductive coding of 30 sources, then applying Sociotechnical Systems (STS) theory as an interpretive lens to reveal the subsystem structure underlying the taxonomy. Unlike Chopra (2025), whose five barrier categories are derived from market reports without theoretical anchoring, and Li et al. (2025), who focus on organizational barriers excluding technical and economic dimensions, this taxonomy: (a) classifies each barrier by agentic specificity – distinguishing barriers unique to agentic AI from those inherited from prior digital transformation waves; (b) maps the inductively-derived dimensions onto STS subsystems based on root-cause analysis, revealing that the social subsystem contains the majority of barriers; (c) analyzes cross-subsystem interactions, revealing self-reinforcing barrier cycles; and (d) identifies the governance calibration problem – balancing control with the autonomy that gives agentic AI its value – as the STS joint optimization challenge that emerges from the interaction analysis, rather than as a barrier within the taxonomy itself.

The article is structured as follows. Section 2 reviews the literature on AI adoption barriers, digital transformation obstacles, and sociotechnical systems theory, organized by analytical tensions rather than themes. Section 3 states the research aims and objectives. Section 4 describes the critical narrative review methodology. Section 5 presents the barrier taxonomy, agentic specificity, and STS subsystem analysis. Section 6 discusses cross-dimensional interactions, counter-evidence, and practical implications. Section 7 offers conclusions and future research directions.

Analysis of recent research and publications. The answer to whether the failures of AI-driven enterprise transformation are primarily due to technical or organizational factors will vary depending on who is asked.

Industry surveys consistently foreground technical infrastructure. Gartner (2025) predicts that over 40% of agentic AI projects will be canceled by 2027 due to escalating costs, unclear business value, or inadequate risk controls – with legacy system incompatibility identified as a key contributing factor. Deloitte's (2026) enterprise survey reports that 48% of organizations cite data searchability and 47% cite data reusability as obstacles to their AI automation strategy. The World Economic Forum (2025) identifies infrastructure constraints as one of three core obstacles to the adoption of agentic AI, alongside trust deficits and data gaps. This evidence indicates that technical readiness is the primary bottleneck.

Academic research offers a different diagnosis. Makarius et al. (2020) argue, based on sociotechnical analysis, that AI implementation carried out without structured employee socialization – regardless of technical capability – is unlikely to create organizational value. Their framework identifies the interaction between AI novelty (how unfamiliar the system is) and AI scope (how broadly it operates) as the critical determinant of integration difficulty, placing organizational adaptation rather than technical capability at the center of the challenge. Papagiannidis et al. (2022), through a comparative case study of three energy-sector firms (15 semi-structured interviews), found that organizations achieving successful AI outcomes did so primarily through organizational practices: standardized processes, employee training programs that addressed both capabilities and job security concerns, and transparent cross-departmental communication structures. Technical factors mattered, but organizational readiness determined whether technical investments translated into value.

This divergence is not a contradiction but a perspective effect. CIO and CTO surveys, which dominate industry reports, naturally emphasize infrastructure because that is their domain of responsibility. Management researchers study organizational dynamics and find organizational barriers because that is where they look. Through the STS lens proposed by Bostrom and Heinen (1977), this false dichotomy dissolves: information system failures stem from the misalignment between social and technical subsystems, not from deficiencies in either subsystem alone. Makarius et al. (2020) explicitly apply this sociotechnical perspective to AI integration, demonstrating that the STS framework can reveal dynamics invisible to single-subsystem analyses. The barrier is not inadequate technology or inadequate organization, but the gap between what the technical system requires and what the organizational system can absorb.

AI governance presents a structural tension that intensifies in agentic contexts. The evidence supporting governance as a transformation enabler is substantial. Papagiannidis et al. (2022) documented that organizations with structured governance across procedural, relational, and structural dimensions reported measurable cost reductions and improved confidence in AI-generated decisions. Hassan et al. (2024), in a scoping review synthesizing 50 studies on AI adoption in healthcare, conclude that a governance structure "can be a key facilitator in ensuring all the elements identified as barriers are addressed appropriately." These findings establish that governance investment correlates with transformation success.

Yet the same governance mechanisms that enable conventional AI can constrain agentic AI. Agentic systems derive their distinctive value from autonomous operation – making decisions, expanding their operational scope, and coordinating across agent networks without continuous human direction (Sapkota et al., 2025). Governance, by definition, constrains this autonomy. This suggests a paradox that the present study examines further in Section 5: organizations with robust governance frameworks may find it easier to restrict agent autonomy than to calibrate it, defaulting to deterministic workflows that negate the value proposition of genuinely autonomous agents. Deloitte (2026) documents a symptom of this dynamic – "agent washing," in which vendors rebrand existing automation capabilities as agentic AI, and organizations misapply agents to cases better suited to simpler tools. While Deloitte attributes agent washing primarily to market hype, overly restrictive governance may compound the problem by making genuinely autonomous deployment organizationally impractical.

Madanchian and Taherdoost (2025), in a review of ethical theories and governance models across multiple jurisdictions, identify a "translational gap" – the disconnect between ethical AI principles developed in academic literature and their actual use in organizational practice. Their analysis highlights significant differences across governance traditions (Anglo-American, Continental European, Asian), suggesting that even standardized governance approaches remain elusive. This regulatory fragmentation compounds the governance calibration challenge: organizations must not only determine the right governance intensity for their context but must do so against a moving regulatory backdrop with no settled standards.

The governance paradox is not a binary choice between control and freedom. It is a calibration problem – the optimal governance intensity varies by decision criticality, organizational risk tolerance, and agent capability maturity. Kasirzadeh and Gabriel (2025) formalize this insight through a four-dimensional agent characterization framework (autonomy, efficacy, goal complexity, generality), demonstrating that governance challenges shift qualitatively across different classes of AI agents – yet no existing framework translates these dimensions into actionable calibration guidance, leaving organizations to oscillate between under-governance (creating risk and trust deficits) and over-governance (suppressing the autonomy that justifies the investment).

The human dimension of AI transformation is commonly framed as "resistance" – employees opposing AI because they fear job displacement. The evidence complicates this framing substantially.

Li et al. (2025), drawing on a survey of over 100 C-suite executives and two dozen cross-industry interviews, identify three barrier domains: people-related barriers (fear of replacement, self-image concerns – workers concealing AI use to avoid appearing incompetent), process-related barriers (treating AI as an overlay rather than workflow transformation), and political barriers (resource hoarding – programmers being 16–18% less likely to recommend AI access to teammates). These findings suggest that what appears as resistance often reflects structural and political dynamics within organizations rather than individual opposition to AI itself.

However, characterizing the human dimension as opposition to AI misreads the evidence. Wiley's (2025) survey of workplace change finds that 68% of employees report feeling excited or curious about AI, and 44% already use it at least weekly. The same survey reveals the actual barriers: 75% lack confidence in using AI effectively, only 18% receive comprehensive change training, and only 34% of managers feel equipped to support AI integration. These statistics describe an enablement failure, not a resistance problem.

The deeper issue is the capacity to absorb change. The average employee experienced 10 planned enterprise changes in 2022, up from just 2 in 2016 (Gartner, cited in O Morain & Aykens, 2023). Gallagher's (2025) sector report lists change fatigue among the top five barriers to organizational success, with 44% of internal communicators identifying it as a key challenge. Wiley (2025) describes the compounding dynamics as a "cascade crisis"–employees battered by successive disruptions before recovering from prior ones, with 96% reporting some degree of stress about workplace change. Beyond change fatigue, AI tool use itself imposes cognitive costs: Bedard et al. (2026) found that workers performing high AI oversight expend 14% more mental effort and experience 19% greater information overload, with productivity declining when employees use more than three AI tools simultaneously.

Classical change management assumes bounded transitions between defined states – Kotter's (1996) eight-step model, for instance, presupposes a "change" with a beginning and an endpoint. Agentic AI invalidates this assumption. Because agent capabilities evolve continuously post-deployment and agent scope expands autonomously, there is no stable "after" state toward which the organization transitions. The change is continuous and open-ended, which means the standard change management prescription – communicate the vision, achieve quick wins, anchor the new state – lacks an anchoring endpoint. This structural mismatch between change management models and the continuous nature of agentic AI transformation suggests that change fatigue may be particularly pronounced in agentic contexts, though empirical evidence specific to agentic AI deployments is not yet available.

The governance, human, and organizational challenges reviewed above raise a more fundamental question: whether agentic AI transformation faces qualitatively new barriers or familiar ones in new packaging. The answer determines whether established frameworks suffice or new analytical tools are needed.

The case for continuity draws on established research on digital transformation barriers. Vogelsang et al. (2019), through 46 expert interviews, identified five barrier categories for digital transformation: missing skills, technical barriers, individual barriers, organizational culture, and environmental barriers. A recent meta-analysis of TOE factors driving AI adoption across industries (Pinto et al., 2025) confirms that seven of eight technology-organization-environment factors that predicted adoption for prior technologies remain statistically significant for AI. From this perspective, agentic AI faces the same barriers as any digital transformation – infrastructure readiness, organizational culture, skills gaps, regulatory environment – and existing frameworks remain adequate.

The case for qualitative difference rests on the distinctive properties that define agentic AI. Sapkota et al. (2025) formally distinguish agentic AI from conventional AI agents, identifying key differentiators including: multi-agent architecture versus single-entity operation; higher autonomy in complex multi-step processes; inter-agent coordination requiring shared memory infrastructure; cross-domain learning scope; and dynamic task decomposition. These properties create barriers without precedent in prior transformation waves. Error propagation across distributed agent networks – where one agent's incorrect output cascades as input to subsequent agents – does not arise in single-system implementations. Emergent behavior from agent interactions can produce outcomes that no individual agent's logic can explain, making governance architecturally inadequate when designed for predictable systems. Continuous scope expansion invalidates the bounded-transition assumption underpinning classical change management models.

The resolution lies in distinguishing between inherited barriers whose mechanisms change in agentic contexts and genuinely novel barriers that emerge solely from agentic properties. Legacy infrastructure incompatibility, for instance, is inherited from every digital transformation wave – but the specific demand profile of agentic AI (real-time autonomous agent integration requiring modern APIs, modular architectures, and secure identity management) represents a qualitatively more demanding version of this inherited barrier. The governance calibration problem, by contrast, is agentic-specific: no prior technology required organizations to balance control with autonomous capability expansion by the technology itself.

Purpose of the article. This study aims to identify, classify, and critically analyze factors that act as barriers to successful enterprise transformation under the influence of agentic AI.

The specific objectives are:

1. To identify barrier dimensions emerging from the convergence of academic literature and industry research on agentic AI enterprise adoption (2019–2026).
2. To classify identified barriers by agentic specificity – distinguishing barriers unique to agentic AI contexts from those inherited from prior digital transformation waves.
3. To map the identified barriers onto STS social and technical subsystems based on root-cause analysis and analyze cross-subsystem interactions, revealing how barriers in one subsystem amplify or create barriers in the other.

Main body of the article. A critical narrative literature review was selected as the methodology for three reasons. First, the field of agentic AI enterprise transformation is emerging – the concept gained traction in academic and practitioner discourse only after 2022, creating an insufficient peer-reviewed corpus for a systematic review. Second, the study's contribution lies in analytical reinterpretation through a specific theoretical lens (STS theory) rather than exhaustive coverage, which aligns with the critical narrative variant's strengths (Grant & Booth, 2009). Third, narrative reviews are established in management research for emerging topics where rigid systematic protocols would be premature (Baumeister & Leary, 1997).

Sociotechnical Systems (STS) theory, developed at the Tavistock Institute from the 1950s onward (Trist, 1981) and applied to information systems by Bostrom and Heinen (1977), holds that organizational performance depends on the joint optimization of two interdependent subsystems: the social (people, roles, processes, culture) and the technical (tools, infrastructure, data, capabilities). Optimizing one subsystem at the expense of the other produces suboptimal outcomes; failures concentrate at the interface between them.

Sources were identified through web-based academic searches targeting peer-reviewed articles, top-tier conference proceedings, and methodologically transparent industry reports. Eight search strings were constructed combining core concepts: agentic AI, AI adoption barriers, digital transformation failures, AI governance challenges, change management and AI, and sociotechnical theory with AI implementation. Sources were also identified through citation chaining from foundational works.

Selection criteria included: direct relevance to barriers or challenges of AI or digital transformation at the enterprise level; publication between 2019 and 2026 (with exceptions for foundational theory works); peer-reviewed journal articles, top-tier conference papers, or industry reports with documented methodology; and English language. Vendor marketing materials, editorials, and abstract-only publications were excluded.

The final analytical corpus comprised 30 sources: 6 classified as T1 (peer-reviewed empirical), 9 as T2 (peer-reviewed review/conceptual), 11 as T3 (industry reports with documented methodology), and 4 as T4 (practitioner sources). The T1-T2 academic ratio of 50% reflects the field's emerging nature, in which practitioner discourse drives academic research. To mitigate this, T3-T4 sources were used solely for contextual data — adoption statistics, trend indicators, and illustrative cases — while argumentative claims were anchored in T1-T2 evidence.

Analysis proceeded in six stages: (1) structured data extraction from each source (barriers identified, categorization scheme, evidence type, context specificity, evaluation scores); (2) open coding of all individual barriers into a master inventory; (3) axial coding to group barriers into dimensions based on conceptual distinctness, support from 3+ sources, and relevance to enterprise-level transformation; (4) agentic-specificity assessment classifying each barrier as agentic-specific, inherited-but-amplified, or inherited. (5) STS subsystem mapping — each barrier was mapped to the social or technical subsystem based on root-cause analysis, regardless of which dimension it was coded into; and (6) cross-subsystem interaction analysis identifying mechanisms through which barriers propagate across the social-technical boundary;

This review does not claim exhaustive coverage. Sources were purposefully selected according to documented criteria and evaluated using a consistent analytical framework. The search relied on web-accessible academic databases and open-access sources rather than institutional database exports, which may introduce selection bias toward more visible and frequently cited works. The T1-T2 source ratio (50%) falls below the conventional 70% threshold, indicating limited academic research on agentic AI transformation. Due to the topic's recency, the empirical base remains limited, and practitioner evidence was used to supplement peer-reviewed findings where explicitly noted.

The analysis identified 29 distinct barriers classified into five dimensions. Table 1 summarizes the taxonomy with agentic-specificity assessment for each barrier.

Table 1. Five-Dimensional Taxonomy of Barriers to Enterprise Transformation Under Agentic AI with STS Subsystem Mapping (originally developed)

Dimension	Barrier	Agentic Specificity	STS Subsystem	Key Evidence
Technological	Legacy system incompatibility	Amplified	Technical	Papagiannidis et al. (2022); Deloitte (2026)
	Data architecture constraints	Amplified	Technical	Deloitte (2026)
	Error propagation in multi-agent systems	Agentic-specific	Technical	Sapkota et al. (2025)
	Explainability deficits	Amplified	Technical	Sapkota et al. (2025); Hassan et al. (2024)
	Cybersecurity vulnerabilities	Amplified	Technical	KPMG (2026)
	Agent interoperability challenges	Amplified	Technical	Sapkota et al. (2025); Bain (2025)
	Real-time processing infrastructure demands	Amplified	Technical	Gartner (2025); KPMG (2026)

The continuation of the table 1

Dimension	Barrier	Agentic Specificity	STS Subsystem	Key Evidence
Organizational	Governance gaps	Amplified	Social	Papagiannidis et al. (2022); Deloitte (2026)
	Strategy absence	Amplified	Social	Deloitte (2026)
	Organizational silos	Amplified	Social	Li et al. (2025); Pandiri (2025)
	Role ambiguity in human-AI collaboration	Agentic-specific	Social	Makarius et al. (2020); MIT SMR/BCG (2025)
	Pilot-to-scale gap	Amplified	Social	Pandiri (2025); Deloitte (2026)
	Leadership AI literacy deficit	Amplified	Social	MIT SMR/BCG (2025); Wiley (2025)
	Cultural inertia toward autonomous systems	Inherited	Social	Vogelsang et al. (2019); Makarius et al. (2020)
Human	Change fatigue/ cascade crisis	Amplified	Social	Wiley (2025); Gartner (cited in O Morain & Aykens, 2023)
	Fear of replacement	Inherited	Social	Li et al. (2025); Papagiannidis et al. (2022)
	Skills and expertise gaps	Amplified	Social	Wiley (2025)
	Resource hoarding	Amplified	Social	Li et al. (2025)
	Cognitive overload from AI-generated outputs	Amplified	Social	Bedard et al. (2026)
	Loss of professional identity	Inherited	Social	Makarius et al. (2020)
Governance & Regulatory	Regulatory vacuum	Amplified	Social	Madanchian & Taherdoost (2025)
	Accountability diffusion	Agentic-specific	Social	Madanchian & Taherdoost (2025)
	Trust deficit	Amplified	Social	WEF (2025); Hassan et al. (2024)
	Cross-jurisdictional regulatory fragmentation	Amplified	Social	Madanchian & Taherdoost (2025)

End of the table 1

Dimension	Barrier	Agentic Specificity	STS Subsystem	Key Evidence
Economic & Financial	Uncertain ROI	Amplified	Technical	Li et al. (2025); Chopra (2025)
	Hidden total cost of ownership	Amplified	Technical	Papagiannidis et al. (2022)
	Vendor lock-in	Amplified	Technical	Bain (2025)
	Measurement and attribution complexity	Amplified	Technical	Deloitte (2026); Li et al. (2025)
	Opportunity cost of delayed adoption	Inherited	Technical	McKinsey (2025)

Of the 29 identified barriers, three are classified as agentic-specific — they arise exclusively from the properties that distinguish agentic AI from conventional AI and prior digital technologies:

1. Error propagation in multi-agent systems. In single-model AI deployments, errors are bound to the model's output. In multi-agent systems, one agent's incorrect output becomes input for downstream agents, creating cascading failure chains that amplify the original error across the system (Sapkota et al., 2025). This barrier has no precedent in prior transformation waves because prior technologies did not feature autonomous inter-system communication of this kind.

2. Role ambiguity in human-AI collaboration. Prior automation technologies had clear functional boundaries – a machine performed task X, a human performed task Y. Agentic AI systems dynamically adjust their operational scope, taking on tasks previously performed by humans and expanding into adjacent functions without explicit reprogramming (Makarius et al., 2020). This continuous boundary shifting creates sustained role ambiguity that conventional job redesign methods — which assume stable role definitions — cannot address.

3. Accountability diffusion. When a multi-agent system produces a decision, responsibility is distributed across the agent that initiated the process, the agents that contributed data or analysis, the developers who designed each agent, the operators who configured the system, and the users who set the objectives (Madanchian & Taherdoost, 2025). This multilayered attribution problem does not arise for single-system technologies where the chain of responsibility is traceable.

The remaining 26 barriers are classified as inherited but amplified (22) or inherited (4). Inherited-but-amplified barriers exist in prior digital transformation contexts but manifest with greater intensity or through different mechanisms under agentic AI conditions. Resource hoarding, for instance, is classified as inherited-but-amplified: resource contention exists in any distributed organizational system, but Li et al. (2025) found that programmers were 16–18% less likely to recommend AI access to teammates — indicating that agentic AI access amplifies hoarding incentives because it confers competitive advantages that sharing diminishes. Similarly, legacy infrastructure incompatibility is a well-documented barrier to digital transformation (Vogelsang et al., 2019). Still, agentic AI's requirement for real-time, autonomous-agent integration with modern APIs, modular architectures, and secure identity management represents a qualitatively more demanding version of this inherited challenge.

Applying the STS lens, each barrier was mapped to one of two subsystems — technical or social — based on root cause (Table 1, STS Subsystem column). The mapping principle is causal origin: a barrier is assigned to the subsystem where its root cause resides, regardless of where its symptoms manifest. For example, organizational silos manifest technically (incompatible agent frameworks) but originate in social structures (departmental incentives, leadership decisions). The trust deficit appears in the governance and regulatory dimensions,n but its root cause is social — it resides in people's willingness to rely on a system. Conversely, uncertain ROI manifests as a strategic concern but originates in the technical subsystem — the complexity of multi-agent architectures makes value attribution inherently difficult.

The mapping yields 12 barriers in the technical subsystem and 17 in the social subsystem. Notably, the STS mapping does not align one-to-one with the five inductively-derived dimensions: barriers from the organizational dimension are split between subsystems based on root cause, while all governance and regulatory barriers map to the social subsystem, because governance is a human-designed structure, even when it regulates technical systems. This cross-cutting pattern confirms that the STS lens reveals subsystem dynamics that the dimensional taxonomy alone does not capture.

The asymmetry (12 technical vs. 17 social) is itself a finding: it suggests that the barrier landscape for agentic AI transformation is weighted toward social challenges. Organizations may overinvest in technical solutions while underinvesting in the social subsystem where the majority of barriers concentrate — and where governance, the social subsystem's primary mechanism for managing the technical subsystem, must perform the joint optimization that STS theory requires.

Five interaction mechanisms involving 10 of the 29 barriers were identified through analytical synthesis of the evidence (Figure 1). These are the barriers for which the literature documents explicit propagation pathways – causal chains where one barrier's persistence or intensification produces or amplifies another. The proposed interactions represent plausible causal pathways supported by converging evidence from multiple sources, though they have not been empirically tested as integrated models. Three pathways (1, 2, 4) cross the boundary between subsystems; two (3, 5) operate within the social subsystem.

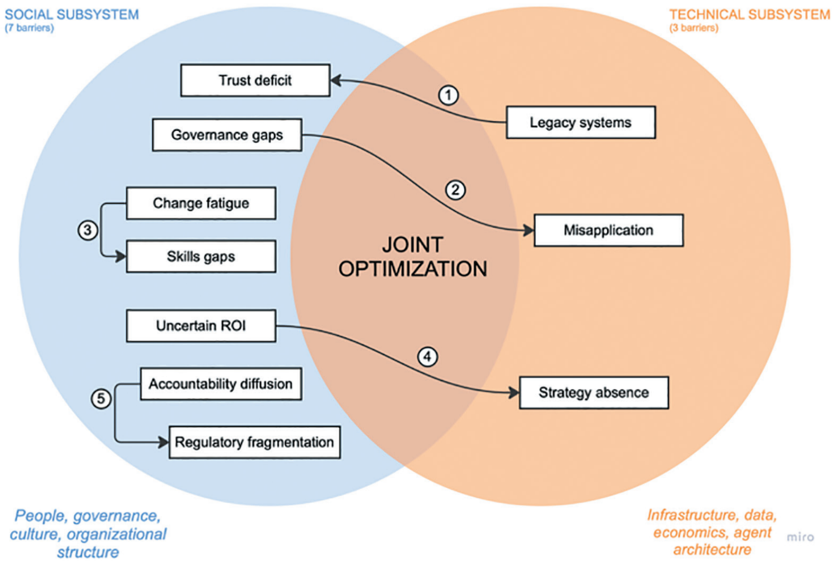

Figure 1. Cross-Subsystem Barrier Interactions Through the STS Lens (originally developed)

1. Legacy systems -> trust deficit (technical -> social). When legacy infrastructure cannot support formal governance mechanisms – automated audit trails, real-time monitoring, access controls – organizations create informal workarounds: manual oversight, spreadsheet-based tracking, ad-hoc approval chains. Papagiannidis et al. (2022) documented this pattern: one studied firm initially relied on insecure Excel file exchanges before implementing API-based data governance. These workarounds are neither scalable nor transparent, eroding stakeholder trust. The pathway illustrates how a technical limitation (legacy systems) creates a social barrier (a trust deficit) by degrading the governance mechanisms that sit between them.

2. Governance gaps -> agent misapplication (social -> technical). When governance frameworks designed for conventional IT fail to account for AI systems that make independent decisions (Deloitte, 2026), organizations lack criteria for distinguishing cases that require genuinely autonomous agents from those better served by simpler automation. The result is agent misapplication – deploying agentic AI where deterministic tools would suffice, producing poor ROI and reinforcing skepticism about scaling. This pathway runs in the opposite direction from 1: a social barrier (governance gaps) produces a technical barrier (misallocated infrastructure investment), contributing to the pilot-to-scale gap – only 11% of organizations have agentic AI in production despite 38% running pilots (Deloitte, 2026).

3. Change fatigue -> skills gaps -> resistance (intra-social cycle). When employees experience high change saturation, they prioritize voluntary training less. With only 18% receiving comprehensive change training (Wiley, 2025), the skills gap widens. Wider skills gaps increase anxiety about AI competence – 75% report lacking confidence in AI use (Wiley, 2025) – which amplifies avoidance behavior. This self-reinforcing cycle operates entirely within the social subsystem: the employees most in need of AI training are the least likely to seek it.

4. Uncertain ROI -> strategy absence (technical -> social). When AI ROI is uncertain – 45% of executives report returns below expectations (Li et al., 2025) – and 35% of organizations have no formal agentic AI strategy (Deloitte, 2026). While neither source documents the causal link directly, the co-occurrence suggests a plausible pathway: when ROI remains uncertain, organizations hesitate to commit resources to governance infrastructure, treating it as overhead rather than an enabler. The absence, in turn, starves governance of organizational mandate. Without governance, scaling failures confirm ROI pessimism – completing the cycle back into the technical subsystem.

5. Accountability diffusion -> regulatory pressure (intra-social cycle). Madanchian and Taherdoost (2025) identify distributed accountability as "a major issue in AI governance," noting that all stakeholders – firms, developers, designers, users – share responsibility for AI outcomes. The causal chain that follows is this study's analytical inference. When no single party can be held accountable for outcomes in a multi-agent system, stakeholder trust erodes, generating demand for regulatory intervention. Yet regulatory frameworks developed for contexts with identifiable human decision-makers struggle to address distributed agent systems, creating regulatory uncertainty that further inhibits adoption. Like 3, this cycle operates within the social subsystem – accountability structures, trust, and regulation are all social constructs – but it is triggered by a technical property (multi-agent decision distribution).

Joint optimization. The interaction map reveals a pattern that STS theory predicts: the three cross-subsystem pathways (1, 2, 4) all pass through governance. Governance is the social subsystem's primary mechanism for managing the technical subsystem – the instrument through which organizations set boundaries on agent autonomy, establish accountability, and allocate resources. When governance is weak, the social subsystem cannot regulate the technical subsystem (2 agent misapplication). When the technical subsystem is inadequate, governance cannot function properly (1 workaround). When economic signals from the technical subsystem are ambiguous, governance is starved of investment (4 underinvestment). This convergence on governance reveals a governance calibration problem that is not itself a barrier but the joint optimization challenge underlying the barrier system: determining how much the social subsystem should constrain the technical subsystem without destroying the autonomous value that makes agentic AI worth deploying. Kasirzadeh and Gabriel (2025) formalize this tension through a four-dimensional agent characterization framework (autonomy, efficacy, goal complexity, generality), demonstrating that governance requirements shift qualitatively as agent autonomy increases – yet no existing framework translates these dimensions into actionable calibration guidance.

The five pathways account for 10 of the 29 identified barriers. Interaction mechanisms for the remaining 19 barriers were not identified in the current evidence base. This does not imply that the remaining barriers operate in isolation – it reflects the limited empirical research on agentic AI transformation to date. As the field matures and longitudinal case studies become available, additional cross-subsystem pathways are likely to emerge, particularly involving barriers such as cybersecurity vulnerabilities, vendor lock-in, and cognitive overload, whose root causes and symptoms span both subsystems.

The barrier taxonomy reveals three findings with implications for how organizations and researchers approach agentic AI transformation.

First, the false dichotomy between technical and organizational barriers dissolves when the inductively-derived taxonomy is mapped onto STS subsystems. The five barrier dimensions reduce to two subsystems – technical (12 barriers: technological and economic dimensions) and social (17 barriers: human, organizational, and governance dimensions). This asymmetry is itself a finding: the barrier landscape is weighted toward the social subsystem, yet organizations disproportionately invest in technical solutions. The most consequential barriers emerge not from either subsystem alone but from misalignment between what the technical subsystem requires and what the social subsystem can absorb. Organizations that invest in infrastructure without parallel investment in governance, skills, and organizational redesign experience the "missing middle" that Pandiri (2025) describes – technically functional pilots that cannot scale because the social subsystem was not co-developed with the technology.

Second, the agentic-specificity analysis indicates that while most barriers (26 of 29) are inherited from prior transformation waves, they manifest with qualitatively different intensity and through different mechanisms in agentic contexts. The implication for practitioners: established digital transformation playbooks address the inherited components of the challenge but miss the agentic-specific dimensions (error propagation, role ambiguity, accountability diffusion) that require novel management approaches.

Third, the cross-subsystem interaction analysis reveals that barriers propagate across the social-technical boundary. Three of the five identified pathways (1, 2, 4) cross the subsystem boundary; two (3, 5) operate as intra-social amplification cycles. Governance – a social subsystem structure – serves as the primary transmission mechanism for cross-subsystem interactions: technical constraints produce governance workarounds (1), governance gaps produce agent misapplication (2), and economic uncertainty starves governance of resources (4). This pattern explains a persistent puzzle in the adoption data: why do barriers persist even when organizations invest substantially in addressing them? The STS answer is that single-subsystem interventions – upgrading infrastructure (technical) or training employees (social) – are undermined by untreated dynamics in the other subsystem. Effective transformation requires simultaneous joint optimization across both subsystems.

The STS analysis yields a further insight: the governance calibration problem that emerges from the interaction analysis is not itself a barrier, but the joint optimization challenge that STS theory predicts will arise at the boundary between social and technical subsystems. STS holds that organizational performance depends on jointly optimizing both subsystems (Bostrom & Heinen, 1977). For agentic AI, this joint optimization materializes as governance calibration — determining how much the social subsystem (rules, oversight, accountability structures) should constrain the technical subsystem (autonomous agents) without destroying the value that autonomy provides. Kasirzadeh and Gabriel (2025) formalize this tension through a four-dimensional agent characterization framework, demonstrating that governance requirements shift qualitatively across different classes of AI agents as autonomy increases. The calibration problem arises because governance must simultaneously enable and constrain. Unlike prior technologies, where more governance was generally better, agentic AI requires finding a calibration point where governance is sufficient to maintain accountability without being so restrictive as to negate the autonomy that distinguishes agentic AI from conventional automation. This is not a solvable optimization problem with a stable solution — the optimal calibration shifts as agent capabilities evolve, organizational experience accumulates, and regulatory environments change.

The taxonomy extends prior barrier research in three directions. Compared to the TOE framework (Tornatzky & Fleischer, 1990) and its AI-specific extensions, the taxonomy adds dimensions that TOE does not capture — the human dimension, distinct from the organizational, and governance, distinct from the environmental. More fundamentally, TOE's additive structure (technology + organization + environment) does not model interaction effects; the STS-informed approach maps the inductively derived dimensions onto subsystems and traces how barriers propagate across the social-technical boundary.

Compared to Chopra (2025), whose five categories (technical infrastructure, organizational design, financial investment, human factors, security and compliance) overlap substantially with the taxonomy's five dimensions, the present work adds: (a) STS subsystem mapping that reveals root-cause dynamics invisible in the dimensional view — including the social subsystem's disproportionate barrier weight, (b) agentic-specificity classification absent from Chopra's generic treatment, and (c) cross-subsystem interaction analysis that Chopra's siloed categories do not permit.

Compared to Vogelsang et al. (2019), whose empirically derived digital transformation taxonomy (skills, technical, individual, organizational culture, environmental) provides the strongest methodological precedent, the present work narrows the context (agentic AI rather than generic digital transformation) while widening the analytical lens (cross-dimensional interactions rather than independent categories).

The analysis must acknowledge three important counter-observations. First, barriers are surmountable. Papagiannidis et al. (2022) documented measurable cost reductions in organizations with effective AI governance. The DBS Bank case (cited in Li et al., 2025) generated $274 million in value through governance frameworks. These success cases do not invalidate the barrier taxonomy; they demonstrate that barriers yield to structured, governance-led interventions, which is itself a finding.

Second, workforce resistance is more accurately characterized as an enablement failure. The 68% excitement rate toward AI (Wiley, 2025) suggests that employees are willing participants when adequately supported. The barrier is not employee opposition but organizational failure to provide training, guidelines, and psychological safety for AI adoption.

Third, the limited T1-T2 academic evidence (50% of the corpus) constrains the strength of argumentative claims. The taxonomy should be understood as an analytically derived framework requiring empirical validation, not as a definitive causal model. The barrier inventory and the proposed interaction mechanisms constitute hypotheses that future empirical research should test.

The taxonomy offers practitioners three actionable insights. First, the agentic-specificity classification helps organizations triage their challenges: inherited barriers (e.g., skills gaps, infrastructure readiness) can be addressed using established digital transformation practices, while agentic-specific barriers (e.g., error propagation, accountability diffusion) require novel approaches. Second, the cross-dimensional interaction map suggests that sequencing matters: addressing governance gaps early may prevent the cascading effects (silos, trust erosion, strategy absence) that governance deficits generate downstream. Third, the change fatigue finding suggests that organizations should assess their workforce's capacity to absorb change before layering agentic AI deployment onto existing transformation programs.

This study has several limitations. The critical narrative review methodology, while appropriate for an emerging field, does not claim exhaustive coverage. The 50% T1-T2 source ratio reflects the limited academic corpus; as the field matures, empirical validation will strengthen or revise the taxonomy. The evidence base is predominantly from North American and European contexts; barriers in emerging economies or cultures with different technology adoption patterns may differ. The cross-dimensional interaction mechanisms are analytically derived — they describe plausible causal pathways supported by evidence from individual studies but have not been empirically tested as integrated models. Finally, the rapid evolution of agentic AI technology means that some barriers identified in 2024–2026 may shift as technology matures, regulations develop, and organizational experience accumulates.

Conclusions and proposals. This study identified, classified, and critically analyzed 29 barriers to enterprise transformation under agentic AI, organized them into a five-dimensional taxonomy (technological, organizational, human, governance and regulatory, and economic and financial), and assessed their agentic specificity.

The study advances the understanding of AI transformation barriers in three ways: (1) it introduces the agentic-specificity distinction, showing that three barriers are unique to agentic AI contexts while 26 are inherited but amplified; (2) it maps the inductively-derived five-dimensional taxonomy onto STS subsystems based on root-cause analysis, revealing that 17 of 29 barriers originate in the social subsystem – a finding that challenges the technical-first framing dominant in industry discourse; (3) it identifies the governance calibration problem – emerging from the cross-subsystem interaction analysis rather than from the barrier taxonomy itself – as the STS joint optimization challenge for agentic AI: governance, as the social subsystem's mechanism for managing the technical subsystem, must simultaneously enable and constrain autonomous operation (Kasirzadeh & Gabriel, 2025).

The taxonomy provides a diagnostic framework for organizations planning or executing agentic AI transformation. The agentic-specificity classification enables practitioners to distinguish challenges that are addressable with established methods from those that require novel approaches. The STS subsystem mapping warns against single-subsystem interventions. It suggests that governance-led sequencing – investing in the social subsystem's capacity to manage the technical subsystem – may prevent downstream cascading failures.

Three specific empirical questions require investigation: (1) whether the proposed cross-dimensional interaction mechanisms operate as described – longitudinal case studies tracking barrier evolution during agentic AI deployment would provide the needed evidence; (2) whether the relative weight of barrier dimensions varies across industries, organization sizes, and cultural contexts – comparative survey research across geographies and sectors would address this; (3) how governance calibration intensity should be determined – action research within organizations deploying agentic AI could develop and test graduated governance models that balance control with autonomous capability.

REFERENCES

1. Bain & Company. (2025). State of the art of agentic AI transformation. Technology Report. https://www.bain.com/insights/state-of-the-art-of-agentic-ai-transformation-technology-report-2025/.

2. Baumeister, R. F., & Leary, M. R. (1997). Writing narrative literature reviews. *Review of General Psychology*, 1(3), 311–320. https://doi.org/10.1037/1089-2680.1.3.311.

3. Bedard, J., Kropp, M., Hsu, M., Karaman, O. T., Hawes, J., & Kellerman, G. R. (2026). Using AI leads to "brain fry." *Harvard Business Review*. https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry.

4. Bostrom, R. P., & Heinen, J. S. (1977). MIS problems and failures: A socio-technical perspective. Part I: The causes. MIS Quarterly, 1(3), 17–32. https://doi.org/10.2307/248710.

5. Chopra, A. (2025). Adoption of AI and agentic systems: Value, challenges, and pathways. California Management Review Insights. https://cmr.berkeley.edu/2025/08/adoption-of-ai-and-agentic-systems-value-challenges-and-pathways/.

6. Deloitte. (2026). Agentic AI strategy. Tech Trends 2026. https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/agentic-ai-strategy.html.

7. Gallagher. (2025). State of the Sector 2024/25. Institute of Internal Communication. https://www.ioic.org.uk/resource/state-of-the-sector-2024-25-change-fatigue-ai-uncertainty-and-leadership-define-internal-communications.html.

8. Gartner. (2025). Predicts 2025: Agentic AI — the evolution of experience. https://www.gartner.com/en/documents/6149877.

9. Grant, M. J., & Booth, A. (2009). A typology of reviews: An analysis of 14 review types and associated methodologies. *Health Information & Libraries Journal*, 26(2), 91–108. https://doi.org/10.1111/j.1471-1842.2009.00848.x.

10. Hassan, M., Kushniruk, A., & Borycki, E. (2024). Barriers to and facilitators of artificial intelligence adoption in health care: Scoping review. *JMIR Human Factors*, 12, e48633. https://doi.org/10.2196/48633.

11. Kasirzadeh, A., & Gabriel, I. (2025). Characterizing AI agents for alignment and governance. *arXiv preprint,* arXiv:2504.21848. https://arxiv.org/abs/2504.21848.

12. Kotter, J. P. (1996). Leading change. Harvard Business Press.

13. KPMG. (2026). AI at scale: How 2025 set the stage for agent-driven enterprise reinvention in 2026. AI Pulse Report. https://kpmg.com/us/en/media/news/q4-ai-pulse.html.

14. Li, J., Zhu, F., & Hua, P. (2025, November). Overcoming the organizational barriers to AI adoption. Harvard Business Review. https://hbr.org/2025/11/overcoming-the-organizational-barriers-to-ai-adoption.

15. Madanchian, M., & Taherdoost, H. (2025). Ethical theories, governance models, and strategic frameworks for responsible AI adoption and organizational success. Frontiers in Artificial Intelligence, 8, 1619029. https://doi.org/10.3389/frai.2025.1619029.

16. Makarius, E. E., Mukherjee, D., Fox, J. D., & Fox, A. K. (2020). Rising with the machines: A sociotechnical framework for bringing artificial intelligence into the

organization. *Journal of Business Research*, 120, 262–273. https://doi.org/10.1016/j.jbusres.2020.07.045.

17. McKinsey & Company. (2025). The agentic organization: Contours of the next paradigm for the AI era. https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization-contours-of-the-next-paradigm-for-the-ai-era.

18. MIT Sloan Management Review & BCG. (2025). The emerging agentic enterprise: How leaders must navigate a new age of AI. https://sloanreview.mit.edu/projects/the-emerging-agentic-enterprise-how-leaders-must-navigate-a-new-age-of-ai/.

19. O Morain, C., & Aykens, P. (2023, May). Employees are losing patience with change initiatives. Harvard Business Review. https://hbr.org/2023/05/employees-are-losing-patience-with-change-initiatives.

20. Pandiri, S. (2025). Bridging the gaps in AI transformation: An evidence-based framework for scalable adoption. California Management Review Insights. https://cmr.berkeley.edu/2025/11/bridging-the-gaps-in-ai-transformation-an-evidence-based-framework-for-scalable-adoption/.

21. Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2022). Toward AI governance: Identifying best practices, potential barriers and outcomes. *Information Systems Frontiers*, 25, 123–141. https://doi.org/10.1007/s10796-022-10251-y.

22. Pinto, A. S., Abreu, A., Pérez Cota, M., Gonçalves, A. J., & Reis, J. L. (2025). A meta-analysis of TOE factors driving organizational adoption of artificial intelligence across industries. *Discover Artificial Intelligence*, 5, 47. https://doi.org/10.1007/s44163-025-00747-2.

23. Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI Agents vs. Agentic AI: A conceptual taxonomy, applications, and challenges. arXiv preprint, arXiv:2505.10468. https://arxiv.org/abs/2505.10468.

24. Tornatzky, L. G., & Fleischer, M. (1990). The processes of technological innovation. Lexington Books.

25. Trist, E. (1981). The evolution of socio-technical systems: A conceptual framework and an action research program–occasional Paper No. 2. Ontario Quality of Working Life Centre.

26. Vogelsang, K., Liere-Netheler, K., Packmohr, S., & Hoppe, U. (2019). A taxonomy of barriers to digital transformation. *14th International Conference on Wirtschaftsinformatik* (WI 2019), Siegen, Germany.

27. Wiley. (2025). New survey reveals growing risk of change fatigue and a "cascade crisis" for employees navigating AI. https://newsroom.wiley.com/press-releases/press-release-details/2025/New-Survey-Reveals-Growing-Risk-of-Change-Fatigue-Cascade-Crisis-for-Employees-Navigating-AI/.

28. World Economic Forum. (2025). Agentic AI: Overcoming 3 obstacles to adoption and innovation. https://www.weforum.org/stories/2025/12/3-obstacles-to-ai-adoption-and-innovation-and-how-to-overcome-them/.