

From Control to Collaboration: Organisational Structures That Enable Successful AI Adoption by Structured Scoping Review

Minwoo Kim^{1*}, James Suh²

¹High School Division, Tecnológico de Monterrey, Santa Catarina, Mexico

²Pinnacle Education, London, United Kingdom

Email: *A01726037@tec.mx

How to cite this paper: Kim, M., & Suh, J. (2026). From Control to Collaboration: Organisational Structures That Enable Successful AI Adoption by Structured Scoping Review. *Open Journal of Business and Management*, 14, 1128-1147. <https://doi.org/10.4236/ojbm.2026.142065>

Received: February 13, 2026

Accepted: March 16, 2026

Published: March 19, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Organisational adoption of artificial intelligence (AI) is increasingly linked to productivity gains and competitive advantage, yet many firms struggle to convert pilots into sustained, organisation-wide value. This structured review synthesises evidence from a focused set of recent peer-reviewed studies examining human-AI decision support, cognitive biases arising from AI recommendations, human-AI collaboration, and AI governance. Three cross-cutting themes recur across the literature. First, research has shifted from emphasising organisational control over AI systems towards designing interdependent human-AI collaboration at scale, in which tasks are allocated according to complementary strengths. Second, individual-level misuse of AI (for example, overreliance, anchoring, or automation bias) can diffuse through group processes and escalate into organisation-wide decision failures. Third, regulatory fragmentation and uncertainty impose compliance and operating-model burdens that can delay adoption and constrain deployment. Building on these themes, the paper proposes practical, structure-sensitive recommendations for six common organisational forms (hierarchical, matrix, flat, hub-and-spoke network, divisional, and team-based). For each structure, recommendations are organised around three implementation lenses: regulatory barriers to adoption, post-adoption risk controls, and organisational conditions that enable effective and productive AI use. The resulting framework is intended to support entrepreneurs and managers in selecting feasible structural interventions aligned with organisational constraints, risk tolerance, and governance capacity.

Keywords

Artificial Intelligence Adoption, Organisational Structure, AI Governance, Human-AI Collaboration, Automation Bias

1. Introduction

Artificial intelligence (AI) has shifted from a specialised analytical capability to a general-purpose organisational technology, encompassing predictive analytics, automation, and increasingly, generative AI and agentic decision support. Many organisations therefore face strategic pressure to adopt AI not only to improve productivity but also to maintain competitiveness in innovation, customer experience, and operational efficiency. However, evidence suggests that implementation frequently stalls after early experimentation. For example, Gartner has projected that a substantial share of generative AI projects will be abandoned after proof of concept when organisations confront poor data quality, inadequate risk controls, escalating costs, or unclear business value (Gartner, 2024). Similarly, large-scale survey evidence indicates that while organisational use of AI is widespread, many firms remain in relatively early stages of scaling and capturing enterprise-level value (McKinsey & Company, 2025). These patterns highlight a persistent “adoption gap” between technical feasibility and sustainable organisational impact.

A central reason for the adoption gap is that AI implementation is not solely a technical integration task. It is a socio-technical change that reshapes workflows, decision rights, accountability, and employee behaviour. Experimental and organisational research shows that AI can influence judgement and decision-making through predictable cognitive mechanisms. In human-AI decision support, users may display overreliance on AI advice even when contextual cues indicate that scepticism is warranted, generating measurable costs (Klingbeil et al., 2024). AI recommendations can also create anchoring effects that distort managerial evaluations, such as assessments of employee performance, especially when the recommendation is treated as a “default” reference point (Carter & Liu, 2025). These behavioural vulnerabilities are particularly relevant in the context of large language models (LLMs), where fluent outputs can appear credible even when ungrounded; documented concerns include poor calibration of confidence, attribution limitations, and hallucinated content, all of which can undermine decision quality if verification practices and governance are weak (Handler et al., 2024).

Governance and regulation further complicate organisational adoption. AI systems can introduce safety, privacy, fairness, and accountability risks, and regulators increasingly expect organisations to demonstrate active risk management rather than passive compliance. The National Institute of Standards and Technology (NIST) provides a widely used, voluntary framework for managing AI risks across the system lifecycle, emphasising governance, mapping context, measuring performance and harms, and managing risks in deployment and monitoring (National Institute of Standards and Technology, 2023). At the international level, the OECD’s AI principles similarly stress trustworthy and responsible AI, including transparency, robustness, accountability, and human-centred values (Organisation for Economic Co-operation and Development, 2019). Yet, applying govern-

ance expectations in practice can be difficult when legal requirements and enforcement approaches differ across jurisdictions and sectors. Recent work on AI and digital governance notes that firms often face fragmented governance arrangements and evolving expectations, requiring them to coordinate across multiple stakeholders and policy layers (Liao et al., 2025). This regulatory and governance environment can delay adoption, force redesign of systems, or raise the cost of scaling AI beyond pilots.

These challenges imply that the organisational structure of a firm matters for AI adoption outcomes. Structure determines how information flows, how decisions are made, how accountability is assigned, and how quickly an organisation can implement controls and course-correct after deployment. For example, highly centralised structures may enable rapid enforcement of governance controls and standardisation, but can reduce local experimentation and slow the development of effective human-AI collaboration. Conversely, decentralised and highly empowered structures can accelerate experimentation and learning, but may be more exposed to inconsistent practices, uneven capability, and the diffusion of individual-level misuse into system-wide decisions. Practical experience in complex settings also suggests that unclear accountability and trust dynamics can impede safe uptake, even when deployment is technically feasible (Gillner, 2024). Taken together, the literature indicates that firms must solve a three-part adoption problem: creating value through effective human-AI work design, controlling post-adoption risks that threaten decision quality and trust, and maintaining compliance under evolving governance and regulatory constraints.

This review synthesises recent peer-reviewed evidence on human-AI decision-making, human-AI collaboration, and AI governance, and translates recurring themes into implementable organisational recommendations. Rather than proposing a single universal model, the paper develops structure-sensitive guidance for six common organisational forms: hierarchical, matrix, flat, hub-and-spoke network, divisional, and team-based. For each structure, recommendations are organised using three implementation lenses aligned with the recurring issues in the literature: 1) regulatory and governance barriers to adoption, 2) post-adoption risks (including overreliance, anchoring, and governance gaps), and 3) enabling organisational conditions that support effective and productive AI use. The objective is to support entrepreneurs and managers in selecting feasible structural interventions that align with organisational constraints, risk tolerance, and governance capacity.

2. Methodology

This study followed a structured scoping review approach and reported the study selection process using the PRISMA-ScR framework (Tricco et al., 2018). The objective of the review was to synthesise recent evidence relevant to organisational AI adoption, with particular attention to three implementation lenses that repeatedly emerge in the literature: 1) regulatory and governance barriers to adoption,

2) post-adoption risks that can degrade decision quality and trust (e.g., overreliance, automation bias, anchoring), and 3) organisational conditions that enable effective and productive AI use at scale.

2.1. Search Strategy

A targeted search was conducted in June 2025 using ScienceDirect and covered publications from June 2020 to June 2025. The search focused on peer-reviewed literature linking artificial intelligence to organisational or business adoption, governance, and decision-making. Keywords were combined using Boolean operators to capture variations in terminology, including “artificial intelligence”, “AI”, “large language model”, and “generative AI”, alongside organisational terms such as “business”, “firm”, “organisation/organization”, “workplace”, “adoption”, “implementation”, “governance”, and “decision support”. A representative query was: (“artificial intelligence” OR AI OR “generative AI” OR “large language model”) AND (adoption OR implementation OR governance OR “decision support”) AND (organisation OR organization OR business OR firm OR workplace).

The search was executed in ScienceDirect in June 2025 with the following constraints: publication date from June 2020 to June 2025; language: English; document type: research articles and review articles (excluding editorials, notes, and short communications). The Boolean query used was: (“artificial intelligence” OR AI OR “generative AI” OR “large language model” OR LLM) AND (adoption OR implementation OR governance OR “decision support” OR “decision making”) AND (organisation OR organization OR business OR firm OR workplace OR management).

ScienceDirect was selected to provide a consistent, full-text accessible corpus for rapid screening within a focused time window. However, restricting identification to a single platform may under-represent relevant work indexed primarily in other databases and venues, including some information systems and management outlets (e.g., ACM/IEEE-indexed proceedings, Web of Science/Scopus-indexed journals not fully covered by ScienceDirect).

2.2. Eligibility Criteria

Studies were eligible if they were written in English and provided conceptual, empirical, or governance-relevant insights on organisational adoption or use of AI, including human-AI decision-making, workplace impacts, or AI governance frameworks applicable to firms. Studies were excluded if they were primarily technical without organisational implications, focused narrowly on clinical practice outcomes without transferable organisational governance insights, or addressed non-AI technologies (e.g., blockchain) as the central subject. In cases where a study was set in a sector such as healthcare, it was retained only if the findings were framed at a governance, implementation, or organisational level and could reasonably generalise beyond clinical decision-making to organisational AI adoption challenges.

2.3. Screening and Selection Process

Records identified through the database search ($n = 382$) were screened by title and abstract. During screening, 364 records were excluded because they did not address organisational or business adoption, did not provide relevant insight on AI use in organisations, or were out of scope for the review's objectives. Full-text articles were sought for retrieval for the remaining records ($n = 18$), and all 18 were retrieved. Following full-text assessment, seven studies were excluded because they did not focus on AI (e.g., centred on other technologies), or they were restricted to clinical practice contexts without governance or organisational insights that could transfer to business adoption. As a result, 11 studies were included from the database search. In addition, two studies were included through pre-defined targeted inclusion criteria established before theme synthesis: 1) the study directly proposes an organisational AI governance/management framework relevant across sectors, and/or 2) the study provides an AI-driven decision support management perspective that explicitly maps to at least one of the three implementation lenses (governance barriers, post-adoption risks, enabling conditions). Targeted inclusion was applied after full-text eligibility assessment of the retrieved set and before final theme consolidation to reduce post hoc selection (Figure 1).

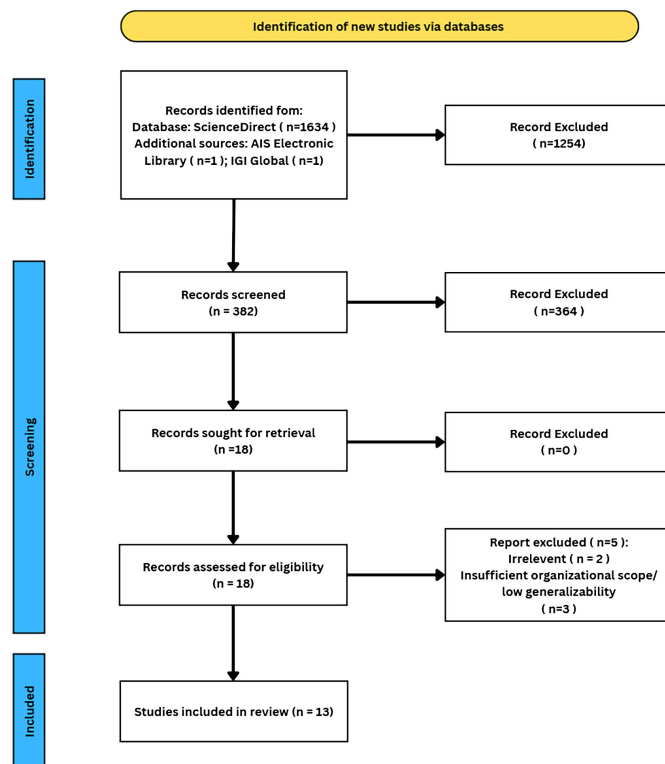


Figure 1. PRISMA-ScR flow diagram of the study selection process. Records identified from ScienceDirect ($n = 382$) were screened by title and abstract, with 364 records excluded. Full texts were assessed for eligibility ($n = 18$); seven articles were excluded after full-text review. Eleven studies were included from the database search, and two additional studies were included through targeted inclusion, yielding 13 included studies in the final synthesis (Tricco et al., 2018).

2.4. Data Extraction and Synthesis

Screening and data extraction were conducted by a single reviewer. To mitigate single-reviewer bias, a verification pass was performed on 1) the full-text exclusion reasons and 2) the mapping of extracted findings to the three trends, with spot-checking of a subset of records and extracted summaries by the co-author/supervisor. For each included study, the reviewer extracted the following information: publication year, study context and domain, study design (e.g., experimental, conceptual, framework-based), the AI technology focus (e.g., decision support systems, diagnostic AI, LLMs), and key findings relevant to organisational adoption. Extracted findings were then synthesised using a thematic approach. Themes were iteratively refined into three cross-cutting trends that recur across the included literature: (1) an evolution from organisational control of AI systems towards designing effective human-AI collaboration; (2) the amplification of individual-level misuse into organisation-wide decision failures through group dynamics and belief diffusion; and (3) regulatory and governance constraints that create barriers to adoption and scaling. These trends were subsequently translated into structure-sensitive recommendations for six organisational forms (hierarchical, matrix, flat, hub-and-spoke network, divisional, and team-based), with each recommendation explicitly organised around regulatory barriers, post-adoption risks, and enabling conditions. Given the focused scope and the heterogeneity of study designs (experimental, conceptual, and framework-based), a formal risk-of-bias appraisal was not applied; instead, the synthesis emphasised recurring mechanisms and governance-relevant implications that converge across multiple study contexts.

3. Results

Thirteen studies met the eligibility criteria and were included in the synthesis (Supplementary data). Collectively, the included literature covered experimental evidence on human reliance and judgement under AI decision support, conceptual and governance frameworks for responsible AI deployment, organisational and workplace implications of human-AI collaboration, and regulatory or governance considerations shaping organisational adoption. The findings were synthesised into three recurring cross-cutting themes (reported below as Trend 1 - Trend 3). These trends were identified through iterative thematic grouping of extracted findings and reflect convergent patterns across multiple study contexts rather than any single domain.

3.1. Study Characteristics (n = 13)

The included evidence comprised: 1) experimental and behavioural studies examining trust, reliance, anchoring, and overreliance under AI advice (Carter & Liu, 2025; Klingbeil et al., 2024), 2) organisational and workplace studies exploring human-AI collaboration outcomes and sensemaking dynamics (Hao et al., 2025; Liu & Li, 2025), 3) conceptual or framework-oriented contributions addressing

decision support, opinion dynamics, and organisational-level diffusion of beliefs (Handler et al., 2024; Wen et al., 2024), and 4) governance and regulation-focused work addressing organisational navigation of AI oversight and policy uncertainty (Gillner, 2024; Liao et al., 2025). Two additional peer-reviewed sources were included to strengthen coverage of organisational AI management and AI-driven decision support in business contexts (Raju et al., 2024; Wang et al., 2024).

3.2. Trend 1: From “Controlling AI” to “Designing Human-AI Collaboration”

Across the included literature, a recurrent pattern was an increasing emphasis on moving beyond a narrow focus on controlling AI systems toward designing interdependent human-AI collaboration that leverages complementary strengths. Governance-oriented work highlights the need for organisational approaches that can manage uncertainty and opacity in AI systems, including “black box” challenges that complicate accountability and assurance (Wang et al., 2024). More recent organisational and workplace research places greater emphasis on how humans and AI jointly perform sensemaking and allocate roles to improve outcomes, suggesting that effective adoption depends on work design and interaction patterns rather than implementation alone (Hao et al., 2025). Similarly, evidence on workplace human-AI collaboration indicates that outcomes depend on context and leadership conditions, reinforcing the need for organisational design choices that shape how AI is used rather than assuming benefits will arise automatically after deployment (Liu & Li, 2025). Overall, this trend supports the interpretation that successful adoption is increasingly framed as an organisational capability in human-AI collaboration rather than a one-off technology rollout (Figure 2).




Figure 2. Schematic of a hierarchical organisational structure. The diagram illustrates a centralised chain of command in which decision authority is concentrated at higher managerial levels and information and accountability flow vertically through successive reporting layers. The figure is used to contextualise how centralisation can support standardised AI governance while shaping escalation pathways for AI-influenced managerial judgement (e.g., anchoring in evaluation decisions).

3.3. Trend 2: Individual-Level AI Misuse Can Amplify into Organisation-Wide Decision Failures

A second robust pattern was the recurring warning that individual-level errors in AI use can scale into broader organisational failures through group processes and decision diffusion. Experimental evidence indicates that people can over-rely on AI advice, incurring costs and increasing the probability of error when human judgement becomes insufficiently critical (Klingbeil et al., 2024). In managerial contexts, AI recommendations can shape judgement through anchoring, potentially biasing performance assessment and other evaluation decisions (Carter & Liu, 2025). Risk-focused analyses in applied settings further demonstrate how automation bias (uncritical acceptance of AI recommendations) can lead to omission and commission errors, emphasising the organisational relevance of individual cognitive vulnerabilities (Moustafa Abdelwanis et al., 2024; Figure 3).




Figure 3. Schematic of a flat (horizontal) organisational structure. The diagram depicts minimal managerial layers and comparatively direct lateral communication across employees, reflecting distributed decision-making and high local autonomy. The figure is used to contextualise how flat structures can accelerate experimentation and learning in human-AI collaboration while increasing exposure to rapid diffusion of individual-level AI misuse into group-level decisions.

The diffusion mechanism is also addressed explicitly in the included conceptual literature. Work on group decision-making and social network opinion dynamics provides a structured account of how beliefs form, spread, and update through interaction, showing how individual judgements can propagate into group-level outcomes (Wen et al., 2024). In parallel, decision-support research on large language models highlights recurring failure modes such as miscalibration and hallucinated content, which can mislead users if outputs are treated as authoritative without verification (Handler et al., 2024). Taken together, these findings indicate that post-adoption risk is not confined to model performance; it also includes predictable behavioural and organisational escalation pathways in which local misuse or overtrust becomes system-level decision error.

3.4. Trend 3: Regulatory Fragmentation and Governance Uncertainty Constrain Organisational Adoption

The third trend was the prominence of governance and regulatory complexity as a practical barrier to adoption and scaling. Governance frameworks emphasise

that organisational AI deployment is shaped by multi-layered questions of accountability, oversight, and alignment between technology and institutional responsibilities (Liao et al., 2025). Sectoral implementation research similarly suggests that uncertainty in accountability, trust, and practical governance can limit real-world uptake even when organisations attempt to implement AI, and that stakeholder expectations and institutional complexity can strongly shape deployment trajectories (Gillner, 2024). In the included business-oriented literature, the governance challenge also appears as a managerial problem of aligning decision support with organisational controls, risk tolerance, and oversight capacity (Raju et al., 2024; Figure 4).

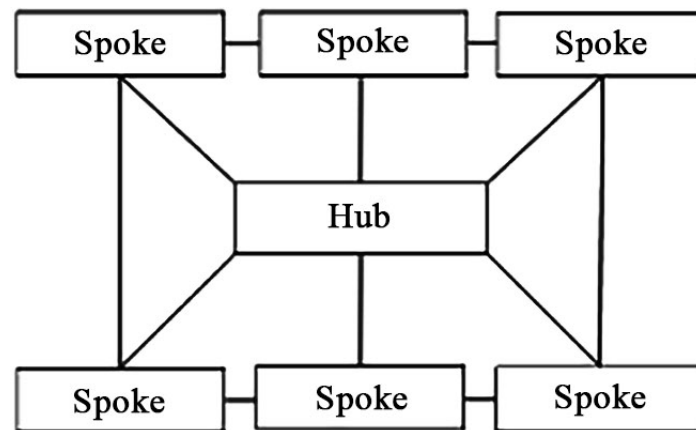


Figure 4. Hub-and-spoke network organisational structure. The diagram illustrates a central hub coordinating multiple spokes (e.g., vendors, suppliers, partners, or subsidiaries), with most inter-spoke coordination mediated via the hub. The figure is used to contextualise governance and accountability challenges in AI adoption across organisational boundaries, including alignment of assurance practices and compliance expectations across partners.

Across these studies, a consistent implication is that adoption strategies must account for governance constraints early, because regulatory uncertainty can impose redesign costs and coordination burdens that reduce feasibility of scaling. This reinforces the rationale for organising subsequent recommendations around regulatory and governance barriers alongside post-adoption risks and enabling conditions for effective use.

4. Discussion

This review set out to clarify why organisational adoption of artificial intelligence (AI) often fails to scale from pilots to reliable, value-generating use, and to translate recurring insights into feasible organisational design recommendations. The synthesis identified three cross-cutting trends: (1) a conceptual shift from “controlling AI” towards designing interdependent human-AI collaboration as an organisational capability (Hao et al., 2025; Liu & Li, 2025); (2) evidence that individual-level misuse and overtrust can amplify into organisation-wide decision

failures via behavioural and diffusion mechanisms (Carter & Liu, 2025; Handler et al., 2024; Klingbeil et al., 2024; Wen et al., 2024); and (3) governance and regulatory complexity as an enduring constraint on adoption feasibility and scaling (Gillner, 2024; Liao et al., 2025; Raju et al., 2024). These trends are mutually reinforcing rather than independent. For example, efforts to increase the intensity of human-AI collaboration (Trend 1) can increase exposure to overreliance and anchoring failures (Trend 2) unless governance and verification mechanisms are strengthened (Trend 3). Consequently, “successful adoption” is best understood as the managed alignment of three conditions: value creation through effective work design, risk containment in post-adoption use, and compliance under uncertainty.

A central implication is that organisational structure matters because it shapes information flows, decision rights, accountability, and the pathways through which AI-generated suggestions become organisational actions. The included literature provides several mechanisms by which structure can moderate adoption outcomes. First, the sensemaking perspective suggests that adoption success depends on how humans and AI coordinate, interpret outputs, and distribute responsibilities, rather than on model capability alone (Hao et al., 2025). Second, behavioural evidence indicates that overtrust and automation bias are not random user errors but systematic cognitive vulnerabilities; where decision processes aggregate or diffuse judgements, such vulnerabilities can become organisation-level failures (Klingbeil et al., 2024; Moustafa Abdelwanis et al., 2024; Wen et al., 2024). Third, governance work indicates that organisations must anticipate oversight, accountability expectations, and operating-model adjustments early, as retrofitting compliance is often costlier and slower than designing governance into workflows from the start (Liao et al., 2025). The discussion below therefore develops structure-sensitive recommendations organised under three implementation lenses: regulatory and governance barriers, post-adoption risks, and enabling conditions for productive use.

4.1. Interpreting Trend 1: Adoption as a Human-AI Collaboration Capability, Not a Deployment Milestone

Trend 1 indicates that the adoption problem has evolved from a question of whether organisations can deploy AI to whether they can design stable human-AI collaboration. This shift is important because it reframes the locus of failure. In many organisations, pilot projects are treated as technical demonstrations, with success measured by model performance metrics or proof-of-concept feasibility. However, the included literature suggests that meaningful adoption is achieved only when work processes are redesigned to allocate tasks appropriately across human and machine strengths and when the organisation develops routines for sensemaking, feedback, and continuous adjustment (Hao et al., 2025). From this viewpoint, “AI adoption” is less analogous to installing software and more analogous to building a new organisational capability that must be trained, monitored,

and iteratively improved.

A critical implication is that organisations should avoid the assumption that “more AI use” is necessarily better. The value of AI depends on task-technology fit and on the integrity of decision processes. For example, in safety-related environments, human-AI collaboration can either improve or hinder performance depending on leadership and contextual dynamism (Liu & Li, 2025). This result challenges simplistic narratives that AI adoption is a monotonic path from manual work to automation. Instead, the adoption pathway is conditional: AI can create value when it is embedded within appropriate controls and aligned with human roles, but it can also degrade performance if embedded in contexts where humans relinquish critical evaluation or where incentives encourage superficial use. Hence, a high standard adoption strategy explicitly defines how collaboration will work (who uses AI, for what tasks, with what verification, and with what escalation procedures) and recognises that “use” should be constrained where risks exceed benefits.

4.2. Interpreting Trend 2: Why Individual Misuse Escalates and How Organisations Can Contain It

Trend 2 is particularly consequential because it explains why otherwise competent organisations can experience sudden decision failures after AI deployment. The included evidence identifies multiple cognitive pathways. One is overreliance: users may treat AI output as authoritative, reducing scrutiny and producing measurable costs even when the system is imperfect (Klingbeil et al., 2024). Another is anchoring: AI-provided starting points can distort subsequent human judgement in evaluation tasks, including performance appraisal (Carter & Liu, 2025). A third is automation bias in applied contexts, which can produce omission and commission errors and thereby create institutional risk (Moustafa Abdelwanis et al., 2024). These vulnerabilities are amplified by LLM-specific failure modes; fluent responses may be accepted as plausible even when they contain fabricated details, and poor calibration may lead users to overestimate reliability (Handler et al., 2024). Importantly, these are not solely “training” problems. Even well-trained staff can succumb to such biases, particularly under time pressure, high workload, or ambiguous accountability.

The diffusion account provides the missing organisational link. Wen et al. (2024) describe how belief formation, influence evaluation, diffusion, and updating produce group-level outcomes. In organisational terms, an individual’s AI-influenced belief can become a team narrative, then a managerial consensus, and finally an institutional decision. This pathway is more likely when information is rapidly shared without robust verification, when teams lack structured dissent, or when the organisation rewards speed over accuracy. Accordingly, mitigation requires more than technical validation: it requires organisational design that reduces uncritical propagation. Practical containment strategies therefore include: 1) verification rituals (e.g., mandatory “evidence checks” for high-impact decisions), 2)

structured dissent or independent review to counter consensus distortion, and 3) decision provenance practices that record how an AI suggestion was used and what human checks were applied. While these measures may appear bureaucratic, the evidence implies that without them, local misuse can become systemic.

4.3. Interpreting Trend 3: Governance Complexity as an Adoption Constraint, Not an Externality

Trend 3 indicates that organisations cannot treat governance and regulation as external constraints to be addressed “after” deployment. Governance shapes the adoption process itself because it affects what can be deployed, where, and under what assurance and accountability expectations. Liao et al. (2025) emphasise that digital governance involves multiple dimensions and questions (captured by their framework) that organisations must address to align technology with institutional responsibilities. Gillner (2024) similarly illustrates how implementation in complex systems is shaped by stakeholder dynamics, uncertainty, and accountability concerns; even when technical capability exists, trust and practical governance can limit uptake. From a business perspective, Raju et al. (2024) underline that AI-driven decision support is a management challenge as well as a technical one, implying the need for operating models that connect AI use to oversight and responsibility.

A critical inference is that governance is not merely a compliance cost; it can be a value enabler when it stabilises trust, clarifies accountability, and reduces the risk of organisational disruption from AI failures. Conversely, weak governance can impose hidden costs: projects are paused after incidents, outputs are ignored because they are not trusted, or adoption becomes fragmented across teams because no one can certify “safe use”. Therefore, an “acceptable” adoption strategy should treat governance capacity (policies, roles, monitoring, escalation) as part of the adoption capability itself.

4.4. Structure-Sensitive Recommendations across Six Organisational Forms

The following recommendations are framed as moderate, feasible changes rather than wholesale redesign. In line with the evidence base, each structure is discussed through three lenses: governance barriers, post-adoption risks, and enabling conditions for productive human-AI collaboration. Importantly, these are not prescriptions that any given structure “should” become another structure. Rather, they represent targeted adjustments that can improve adoption readiness while preserving the advantages that motivate the structure in the first place.

4.4.1. Hierarchical Structures

Hierarchical structures often concentrate decision rights, which can support standardisation and governance enforcement. This centralisation can be advantageous under Trend 3 because it allows the organisation to implement consistent rules for model access, acceptable use, and auditability. However, hierarchical decision

processes can also amplify Trend 2 risks if managers treat AI outputs as authoritative shortcuts under time pressure. Anchoring effects are particularly relevant in settings where performance evaluation and resource allocation decisions have high consequences, because an AI recommendation can become the reference point for judgement (Carter & Liu, 2025). Moreover, if lower-level employees are discouraged from challenging managerial conclusions, errors may not be corrected before they scale.

A feasible intervention is to embed structured feedback loops that enable “upward correction” without undermining the chain of command. One practical approach is to introduce an internal “decision provenance” protocol for AI-supported managerial decisions. For example, when AI support is used in evaluation or resource allocation, the decision record would specify what inputs were used, what human checks were applied, and what alternative interpretations were considered. This does not require flattening the hierarchy; it requires making the decision process auditable and contestable. Because LLMs can hallucinate or present ungrounded content with persuasive fluency (Handler et al., 2024), the protocol should mandate verification for high-impact conclusions, such as requiring managers to confirm claims through primary evidence (records, performance metrics, peer feedback) rather than accepting AI summarisation as proof.

To support Trend 1, hierarchical organisations can treat AI as a “collaboration tool” rather than a “manager substitute” by defining where AI is used to expand situational awareness rather than to decide. For instance, AI can be used to synthesise employee feedback at scale (pattern recognition), while humans apply contextual judgement, resolve ambiguities, and decide priorities (Hao et al., 2025). However, the ethical risks noted in your draft (privacy and surveillance) are real and must be treated as governance issues (Trend 3). If employee feedback is collected, the organisation should minimise personal data, limit access, apply retention controls, and separate developmental feedback from disciplinary processes to prevent chilling effects and mistrust, which would undermine adoption through reduced participation and diminished trust (Gillner, 2024).

4.4.2. Matrix Structures

Matrix structures create dual accountability lines, which can increase coordination burdens but also support cross-functional integration. Under Trend 1, matrix structures can be advantageous because they allow AI-enabled ideas to move across functional boundaries, encouraging iterative refinement and shared learning. However, matrix structures may be especially vulnerable to Trend 2 escalation because the same AI output can be interpreted differently by functional and project leaders, potentially generating conflicting directives. Diffusion dynamics are relevant here: belief formation and influence evaluation may occur within each reporting line, leading to distinct group-level outcomes that are difficult to reconcile (Wen et al., 2024). When combined with LLM failure modes (e.g., miscalibration and hallucination), the risk is not only that a team is wrong, but that teams disagree confidently in incompatible directions (Handler et al., 2024).

A governance-oriented solution is to formalise an “AI assurance intermediary” function for high-impact decisions, analogous to a lightweight review role. This intermediary is not intended to police every AI interaction; that would be infeasible and counterproductive. Rather, it is intended to review AI use in decisions that cross reporting lines or materially affect risk exposure (e.g., compliance-sensitive deliverables, safety-related recommendations, high-stakes performance decisions). The intermediary would check whether AI outputs are supported by evidence and whether the decision process includes independent verification steps. This intervention directly targets Trend 2 while preserving the matrix advantage of cross-functional collaboration.

Under Trend 3, matrix structures should define who owns regulatory interpretation and risk acceptance to avoid “accountability gaps”. [Liao et al. \(2025\)](#) imply that governance involves multiple layers of responsibility; matrix organisations therefore benefit from an explicit rule that assigns final risk acceptance to a designated role (e.g., a governance committee or accountable executive) rather than leaving it to negotiation between managers. This does not remove collaboration; it clarifies where escalation ends, thereby reducing paralysis and enabling adoption to proceed with controlled risk.

4.4.3. Flat Structures

Flat structures typically facilitate rapid communication and local experimentation, which can support Trend 1 by enabling teams to iteratively integrate AI into workflows and adapt quickly. However, the same features can intensify Trend 2 because diffusion and consensus formation can occur quickly, and social dynamics may discourage dissent once an AI-supported narrative becomes dominant. Overreliance and anchoring are particularly problematic when there are few hierarchical checkpoints and when speed is rewarded ([Carter & Liu, 2025](#); [Klingbeil et al., 2024](#)). Moreover, LLM outputs can appear coherent even when ungrounded, potentially accelerating belief formation and diffusion ([Handler et al., 2024](#); [Wen et al., 2024](#)).

For flat structures, the adoption priority is to preserve experimentation while ensuring that high-impact decisions include structured scepticism. A feasible mechanism is to introduce “decision tiers” rather than adding managerial layers. Low-impact decisions can remain decentralised and fast, while high-impact decisions (those affecting compliance, safety, major resource allocation, or personnel outcomes) trigger a lightweight independent review or “red team” check. The red team role can rotate to avoid creating a new hierarchy, but it should be defined so that dissent is institutionalised rather than dependent on personality. This directly counteracts diffusion-driven escalation by ensuring that at least one actor has a formal mandate to challenge AI-supported conclusions.

Under Trend 3, flat structures face a specific risk: governance may become fragmented because there is no central authority to enforce consistent controls. [Liao et al. \(2025\)](#) suggest that governance requires alignment of responsibilities across stakeholders; therefore, flat organisations should adopt minimal shared standards

for AI use (acceptable-use rules, data-handling constraints, documentation expectations for high-impact use) while leaving implementation flexibility at the team level. The goal is to achieve “governance coherence without reintroducing bureaucracy”. This is critical because if governance is inconsistent, adoption becomes risky and credibility is undermined after the first failure.

4.4.4. Hub-and-Spoke Network Structures

Hub-and-spoke networks rely on a central coordinating organisation (hub) that integrates contributions from external partners (spokes). This structure can provide flexibility under Trend 3 because partnerships can be reconfigured when regulatory or governance constraints change. Yet, the same networked arrangement can complicate governance and accountability because responsibility for AI outputs may be distributed across organisations. Gillner (2024) illustrates that implementation in complex systems involves trust and coordination issues; in networked collaborations, these issues can be amplified because partners may have different risk tolerances, documentation standards, and verification practices.

A practical recommendation is to standardise “AI interface contracts” between hub and spokes. This means specifying, at minimum, what data can be shared, what model outputs can be used for, what verification steps are required before outputs are acted upon, and how incidents are escalated. This approach operationalises Trend 3 governance within the structure’s natural contract boundaries. For Trend 2, the hub should require that any AI-supported recommendation that affects shared outcomes includes an evidence note describing the basis for the recommendation and any known uncertainties. This is particularly important given LLM risks such as hallucination and miscalibration (Handler et al., 2024). The objective is not to eliminate risk but to prevent uncontrolled diffusion of unsupported claims across partner organisations.

4.4.5. Divisional Structures

Divisional structures often combine local autonomy with central oversight, which can be advantageous for balancing Trends 1 - 3. Divisions can experiment and tailor AI use to context (supporting Trend 1), while headquarters can define governance standards (supporting Trend 3). The key risk is fragmentation: divisions may adopt divergent AI practices that are difficult to reconcile or audit. Under Trend 2, a critical concern is that decision quality may vary across divisions, and diffusion of flawed practices may occur if successful-looking but risky approaches spread informally. Given the evidence that individuals can over-rely on AI and that such reliance produces measurable costs (Klingbeil et al., 2024), divisional organisations should be cautious about allowing AI “best practices” to propagate without evaluation.

A feasible approach is to implement “federated governance”: headquarters defines minimum requirements (documentation for high-impact AI use, verification expectations, incident reporting), while divisions retain flexibility in tooling and workflow design. To support learning without uncontrolled diffusion, divi-

sions can submit short post-implementation reviews of AI use cases, focusing on observed failure modes (e.g., anchoring incidents, overreliance) and mitigation effectiveness. This directly engages Trend 2 by making failure modes visible and actionable rather than hidden. Additionally, divisional structures can use “comparative audits” to identify whether certain divisions systematically exhibit higher error rates or lower trust, which may signal governance gaps. These mechanisms align with the idea that adoption success is an organisational capability rather than a one-time project (Hao et al., 2025).

4.4.6. Team-Based Structures

Team-based structures, particularly those relying on cross-functional teams and rotating membership, can be conducive to Trend 1 because they enable rapid recombination of expertise and iterative experimentation. However, they can also intensify Trend 2 because belief diffusion can be fast when teams share information frequently and rely on peer evaluation. Anchoring effects may be salient in peer-based performance assessment if AI-generated summaries become default narratives about a person’s contribution (Carter & Liu, 2025). Moreover, LLM failure modes can affect team deliberations by introducing ungrounded but persuasive claims (Handler et al., 2024).

A high-standard intervention is to embed “verification and dissent roles” within team routines. Rather than assigning a permanent facilitator to police AI use, teams can designate a rotating “verification lead” for decisions that have external consequences or high risk. The verification lead checks whether AI-provided claims are supported by evidence and whether uncertainties are documented. This approach preserves team autonomy while institutionalising scepticism. To reduce anchoring in evaluation, teams should avoid using AI summaries as primary appraisal evidence; instead, AI can be used to organise raw inputs (notes, outputs, peer feedback), while humans perform the evaluative judgement and explicitly note where AI could bias interpretation (Carter & Liu, 2025). This practice is consistent with the broader evidence that overtrust is costly and that critical evaluation must be preserved (Klingbeil et al., 2024).

4.5. Cross-Cutting “Generalised Characteristics” and Organisational Dilemmas

Beyond structure-specific recommendations, the review supports several cross-cutting dilemmas that organisations must navigate. First, communication and empowerment can support Trend 1 by enabling active human-AI collaboration and iterative learning. However, increased communication also accelerates diffusion pathways, making Trend 2 escalation more likely when verification is weak (Wen et al., 2024). Thus, the key dilemma is not “more versus less communication” but “communication with or without scepticism”. The evidence suggests that adoption is most robust when communication is paired with institutionalised verification and structured dissent (Handler et al., 2024; Klingbeil et al., 2024).

Second, control and flexibility are not simple opposites. Trend 3 implies that

governance is necessary to stabilise adoption under uncertainty, but overly restrictive control can suppress the experimentation needed to discover effective human-AI collaboration patterns (Hao et al., 2025). The appropriate balance depends on context: highly regulated or high-harm domains require stronger controls, whereas low-risk domains can tolerate more flexibility. A practical implication is that organisations should adopt tiered governance: strong controls for high-impact decisions and lighter controls for low-impact tasks. This approach is consistent with the empirical evidence that harms and costs emerge when AI advice is accepted uncritically in consequential contexts (Klingbeil et al., 2024; Moustafa Abdelwanis et al., 2024).

Third, trust is both a prerequisite and a risk. Organisations need sufficient trust for AI to be used, but too much trust can produce overreliance and anchoring failures (Carter & Liu, 2025; Klingbeil et al., 2024). Trust should therefore be calibrated rather than maximised. One implication is that adoption programmes should not aim to “increase trust” in general; they should aim to increase justified trust by connecting outputs to evidence, uncertainty communication, and accountability structures (Handler et al., 2024; Liao et al., 2025). In practice, this means framing AI outputs as decision inputs with known limitations and designing processes that keep human responsibility explicit.

4.6. Critical Appraisal and Limitations of Inference

Several limitations constrain the strength of claims that can be made from this review. First, the search strategy relied on a single database and a targeted set of keywords, which increases the risk of missing relevant studies and may bias the included evidence towards certain disciplines or publishers. Second, while the review synthesised themes across diverse studies, it did not conduct a formal quality appraisal or risk-of-bias assessment; consequently, the synthesis should be interpreted as a structured thematic integration rather than as a quantitative estimate of effect sizes. Third, several included studies are context-specific (e.g., healthcare settings), and although governance and behavioural mechanisms may generalise, the extent of generalisability cannot be assumed without further field validation (Gillner, 2024; Moustafa Abdelwanis et al., 2024). Fourth, the presence of conceptual and framework papers alongside experimental studies means that the evidence base combines different epistemic claims (normative guidance, mechanism proposals, empirical results). The discussion therefore emphasises plausible mechanisms and design implications rather than definitive causal conclusions.

A further limitation concerns the translation from trends to organisational structure prescriptions. Although structure plausibly moderates adoption outcomes by shaping diffusion, accountability, and decision rights, the included evidence does not directly test “structure-specific interventions” in controlled settings. The recommendations offered here should therefore be understood as theory-informed design hypotheses grounded in recurring behavioural and governance mechanisms, rather than as experimentally validated prescriptions. This distinction is

important for responsible scholarship: making strong claims about what each structure “requires” would exceed the evidence base. Instead, the contribution of this review is to provide a coherent lens for organisational diagnosis and intervention selection that is consistent with the mechanisms documented in the included literature (Handler et al., 2024; Hao et al., 2025; Wen et al., 2024).

4.7. Implications for Practice and Future Research

For practitioners, the principal implication is that AI adoption should be planned as an organisational capability-building programme. This implies budgeting not only for technical development but also for governance roles, verification routines, training in cognitive pitfalls, and change management to embed human-AI collaboration in everyday practice. A second implication is that organisations should explicitly identify where they are most vulnerable on the three-lens framework: some may be constrained primarily by governance uncertainty, others by behavioural risk (overreliance), and others by weak enabling conditions (poor work design for collaboration). Diagnosing the dominant constraint can prevent costly investments in solutions that do not address the true bottleneck.

For research, a priority is field validation of structure-sensitive interventions. The literature would benefit from comparative studies testing whether tiered governance, rotating dissent roles, or assurance intermediaries reduce error escalation and improve adoption durability in real organisations. A second priority is measurement: future studies should operationalise “adoption success” beyond deployment metrics, incorporating sustained use, decision quality, incident rates, trust calibration, and compliance outcomes. Finally, future work should investigate ethical and privacy risks in organisational monitoring mechanisms used to manage AI adoption, as interventions such as feedback databases and audit logs can inadvertently create surveillance cultures that undermine trust and participation (Gillner, 2024; Liao et al., 2025).

Overall, the synthesis supports the conclusion that successful organisational AI adoption is a multi-dimensional management problem requiring coordinated attention to collaboration design, behavioural risk containment, and governance under uncertainty. Structural interventions are not sufficient on their own, but structure-aware design can reduce predictable failure modes and increase the probability that AI is used productively, safely, and sustainably.

5. Conclusion

This structured review examined recent evidence on organisational adoption of artificial intelligence (AI) and synthesised three recurring themes that shape whether adoption succeeds beyond pilot implementation. First, the literature increasingly frames adoption as an organisational capability in human-AI collaboration, emphasising work design, sensemaking, and complementary allocation of responsibilities rather than technology deployment alone (Hao et al., 2025; Liu & Li, 2025). Second, empirical and conceptual work consistently shows that predictable cog-

nitive vulnerabilities, including overreliance and anchoring, can cause individual-level misuse of AI to escalate into organisation-wide decision failures through diffusion and group dynamics (Carter & Liu, 2025; Handler et al., 2024; Klingbeil et al., 2024; Wen et al., 2024). Third, governance and regulatory complexity functions as a practical constraint on scaling, shaping operating models, accountability, and trust in deployed systems (Gillner, 2024; Liao et al., 2025).

Based on these themes, the paper proposed structure-sensitive recommendations for six common organisational forms (hierarchical, matrix, flat, hub-and-spoke network, divisional, and team-based). Rather than advocating for comprehensive restructuring, the recommendations focus on feasible, moderate interventions that align with three implementation lenses: regulatory and governance barriers, post-adoption risk controls, and enabling conditions for productive AI use. The central conclusion is that AI adoption success depends on aligning value creation, risk containment, and governance under uncertainty. Organisations that treat adoption as capability-building, supported by calibrated trust, verification routines, and clear accountability, are more likely to sustain benefits while avoiding predictable failure modes.

Acknowledgements

The author gratefully acknowledges Dr James Suh for providing academic supervision and constructive feedback on the study design, manuscript structure, and clarity of argumentation. The author also thanks educators and peers who offered general comments on academic writing and presentation. Any remaining errors or omissions are the sole responsibility of the author.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- Carter, L., & Liu, D. (2025). How Was My Performance? Exploring the Role of Anchoring Bias in AI-Assisted Decision Making. *International Journal of Information Management*, 82, Article 102875. <https://doi.org/10.1016/j.ijinfomgt.2025.102875>
- Gartner (2024). Gartner Predicts 30% of Generative AI Projects Will be Abandoned after Proof of Concept by End of 2025. *Gartner Newsroom*. <https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025>
- Gillner, S. (2024). We're Implementing AI Now, So Why Not Ask Us What to Do? —How AI Providers Perceive and Navigate the Spread of Diagnostic AI in Complex Healthcare Systems. *Social Science & Medicine*, 340, Article 116442. <https://doi.org/10.1016/j.socscimed.2023.116442>
- Handler, A., Larsen, K. R., & Hackathorn, R. (2024). Large Language Models Present New Questions for Decision Support. *International Journal of Information Management*, 79, Article 102811. <https://doi.org/10.1016/j.ijinfomgt.2024.102811>
- Hao, X., Demir, E., & Eyers, D. (2025). Beyond Human-in-the-Loop: Sensemaking between

- Artificial Intelligence and Human Intelligence Collaboration. *Sustainable Futures*, 10, Article 101152. <https://doi.org/10.1016/j.sftr.2025.101152>
- Klingbeil, A., Grützner, C., & Schreck, P. (2024). Trust and Reliance on AI—An Experimental Study on the Extent and Costs of Overreliance on AI. *Computers in Human Behavior*, 160, Article 108352. <https://doi.org/10.1016/j.chb.2024.108352>
- Liao, S. M., Haykel, I., Cheung, K., & Matalon, T. (2025). Navigating the Complexities of AI and Digital Governance: The 5W1H Framework. *Journal of Responsible Technology*, 23, Article 100127. <https://doi.org/10.1016/j.jrt.2025.100127>
- Liu, Y., & Li, Y. (2025). Does Human-AI Collaboration Promote or Hinder Employees' Safety Performance? A Job Demands-Resources Perspective. *Safety Science*, 188, Article 106872. <https://doi.org/10.1016/j.ssci.2025.106872>
- McKinsey & Company (2025). *The State of AI in 2025: Agents, Innovation, and Transformation*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- Moustafa Abdelwanis, H., Hamdan Khalaf Alarafati, H., Saleh, M., & Can, M. (2024). Exploring the Risks of Automation Bias in Healthcare Artificial Intelligence Applications: A Bowtie Analysis. *Journal of Safety Science and Resilience*, 5, 460-469. <https://doi.org/10.1016/j.jnlssr.2024.06.001>
- National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. U.S. Department of Commerce. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- Organization for Economic Co-Operation and Development (2019). *Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449)*. OECD Legal Instruments. <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449>
- Raju, P., Arun, R., Turlapati, V. R., Veeran, L., & Rajesh, S. (2024). Next-Generation Management on Exploring AI-Driven Decision Support in Business. In S. S. Rajest, S. Moccia, & B. Singh (Eds.), *Optimizing Intelligent Systems for Cross-Industry Application* (pp. 61-78). IGI Global. <https://doi.org/10.4018/979-8-3693-8659-0.ch004>
- Tricco, A. C., Lillie, E., Zarin, W., O'Brien, K. K., Colquhoun, H., Levac, D. et al. (2018). PRISMA Extension for Scoping Reviews (Prisma-ScR): Checklist and Explanation. *Annals of Internal Medicine*, 169, 467-473. <https://doi.org/10.7326/m18-0850>
- Wang, B. Y., Boell, S., Li, C. X., & Chen, E. (2024). Responsible Management for Dynamic Black Box AI: A Cybernetic Approach. In D. Vogel, H. Gewalt, A. Sapsomboon, A. Schwarz, C. Cheung, S. Laumer, & J. Thatcher (Eds.), *Proceedings of the 45th International Conference on Information Systems (ICIS 2024): Digital Platforms for Emerging Societies* (Article 11). AIS Electronic Library. https://aisel.aisnet.org/icis2024/it_implement/it_implement/11/
- Wen, T., Zheng, R., Wu, T., Liu, Z., Zhou, M., Syed, T. A. et al. (2024). Formulating Opinion Dynamics from Belief Formation, Diffusion and Updating in Social Network Group Decision-Making: Towards Developing a Holistic Framework. *European Journal of Operational Research*, 325, 381-399. <https://doi.org/10.1016/j.ejor.2024.12.015>