The Organizational Intelligence Loop: A Socio-Technical Framework for Adaptive Enterprise AI Adoption

By: Yumi Willems Kimura
Information and Knowledge Strategy (IKNS)
Columbia University School of Professional Studies
Independent Research Paper
Advisor/Instructor: Katrina Pugh, Ph.D

December 2025

Table of Contents

Abstract

Organizations are rapidly adopting generative AI (GenAI) and agentic AI, yet outcomes vary widely because technical capability often advances faster than the organizational knowledge structures and socio-behavioral systems required to sustain AI use at scale. Drawing on practitioner interviews, systematic analysis of public industry commentary, and practitioner evidence derived from large-scale enterprise behavioral and collaboration metadata from a knowledge-sharing platform, this study examines how trust, informal influence, expertise visibility, knowledge flow, and workflow behavior shape adoption trajectories. Findings suggest that AI model accuracy and access are insufficient: adoption appears stronger when organizational knowledge is structurally governable (clear ownership, consistent definitions, and reliable provenance), when people-side context is observable (expertise signals, trust pathways, and influence bottlenecks), and when learning is reinforced through workflow-embedded feedback loops, rather than episodic training alone. Across sources, “architectural maturity” (reliable, integrated systems and coherent data foundations) emerges as an enabling precondition for these mechanisms, though it is not sufficient to produce sustained adoption on its own.

To address this gap, the study introduces the Organizational Intelligence Loop (OIL), a socio-technical model that integrates people-side knowledge, information governance, workflow telemetry, and adaptive AI feedback into a unified learning architecture. OIL is positioned as an organizational-layer agentic AI design framework, specifying the governed, runtime-retrievable context and intervention pathways that enterprise agents (and humans) require to adapt safely over time. In doing so, the paper articulates two complementary categories of design and adaptive management: (1) Organizational Context & Adoption Infrastructure for Enterprise AI; and (2) Organizational-Layer Agentic AI Design. Limitations include observational constraints, bounds in coverage and representativeness, and the absence of an end-to-end implemented and empirically validated retrieval/grounding stack. Future research priorities include policy-aware contextual grounding, multi-agent oversight, and alignment “drift” in long-running enterprise agents.

Introduction

Artificial intelligence (AI) is no longer optional. Tools capable of generating content, synthesizing knowledge, and reasoning across systems have shifted from pilot experiments to daily operations. Yet adoption remains uneven and polarized (McKinsey & Company, 2024). On one end, some organizations rush headlong into AI integration, deploying models without adequate governance or alignment. On the other, organizations hesitate, constrained by ethical, compliance, or reputational risk. Between these extremes lies a larger group eager to implement AI but unable to realize its promised productivity gains.

In this study, “AI” refers to computational systems that perceive patterns, make inferences, and recommend or execute actions based on data. This includes analytical AI, which analyzes organizational behavior and knowledge flows, and practical tools such as translation or recommendation engines that personalize user experiences. “Generative AI(GenAI)” denotes models that create new outputs such as text, images, or code by learning from existing data. “Agentic” AI extends these capabilities through autonomous or semi-autonomous systems that reason across environments, adapt to feedback, and perform multi-step tasks. Within this context, an “AI agent” is an instantiated unit of Agentic AI that can pursue defined goals, interact with humans or other systems, and orchestrate actions across applications or workflows.

This study concludes that even forward-leaning adopters are leaving value on the table. The gap is not technical capability but organizational readiness—specifically, the absence of an integrated “people layer” that connects workplace human behavioral patterns, knowledge flow, and learning. Research consistently shows that technology underperforms when social and cognitive systems are overlooked (Shao et al., 2024). Without shared understanding, trust, and collaboration, enterprise AI initiatives may similarly plateau, regardless of investment level.

This study focuses on enterprise GenAI adoption, particularly large-scale assistants and agentic workflows that integrate knowledge and decision support into daily operations and rely on conversational or generative interaction. These systems, like Walmart’s *My Assistant*, go beyond point solutions to retrieve, synthesize, and act on cross-system information in real time (Wilkinson, 2023). The emphasis here is on governed, work-embedded use: AI integrated into existing workflows, instrumented for tracking data provenance, safety, and feedback, and capable of learning from people, information, and process signals over time. The underlying models in enterprise GenAI are trained primarily on organizational knowledge curated and governed internally, distinguishing them from products built mainly on external data. External sources may be incorporated, but in general they are reviewed or contextualized by internal experts. Direct inputs from non-employees—including contractors and external consultants—are excluded from this study.

At its core, enterprise AI adoption *may be* a knowledge management (KM) problem, dependent on how knowledge is created, shared, and applied across networks of people and systems. When these networks fragment, AI may struggle to integrate into them effectively. Traditional enterprise KM investments have emphasized repositories, structured data, and codified knowledge, while tacit knowledge (which is revealed socially) remains difficult to capture and has received fewer investments. This study hypothesizes that the inability to manage tacit expertise, trust, and collaboration patterns contributes directly to uneven AI adoption across organizations and teams. Tacit knowledge often travels through informal networks, which are

typically invisible but strongly shape whether employees feel safe experimenting with new tools or resisting change(Edmondson, 2019; Kotter, 1996). Advances in organizational network analysis and behavioral analytics now make these dynamics visible and measurable (Borgatti, Everett, & Johnson, 2018), offering a way to build the missing infrastructure for sustainable adoption.

The central question guiding this research is how organizations can build the socio-technical conditions required for adaptive and sustainable AI adoption. In this context, “People-Side KM” refers to the human and social dimensions of KM, including the trust, collaboration, and learning behaviors that shape how knowledge is created, shared, and used across an organization. (Nonaka & Takeuchi, 1995; Cross, Parker, Prusak, & Borgatti, 2001). In this study, people-side KM is treated as a designable infrastructure rather than an emergent by-product of collaboration.

In this framing, semantic/metadata work (taxonomies, ontologies, and provenance standards) sits within the curation-and-governance foundation (RQ1), alongside the enabling conditions that make those standards operational—such as identity/access controls, system integration, and reliable data pipelines. People-side KM (operationalized through ONA) addresses the relational and behavioral infrastructure that shapes adoption and learning (RQ2–RQ3). Together, these foundations define the organizational context and guardrails within which enterprise AI systems can learn, route, and act reliably.

This study explores three interrelated questions:

- **Curation and Governance**
  How can information architecture, metadata standards, and governance practices—including semantic/metadata normalization where relevant—be designed, socialized, and measured to improve knowledge transparency, provenance, and trustworthiness in AI-supported environments? This includes examining how baseline “architectural maturity” (coherent integration across core collaboration, workflow, and knowledge systems; stable identities and access controls; and reliable data pipelines) enables a consistent “single source of truth” through master data management and curation workflows, and how organization-specific context (e.g., org structure, decision rights, and responses to prior interventions) can be governed as a first-class knowledge asset that agentic systems must rely on safely.

- **Behavioral Insights**
  How can behavioral and network data identify defensive routines, change influencers, knowledge bottlenecks, and trust patterns, and how can evidence-based nudges strengthen literacy and confidence in AI use?

- **Systems-Level Learning Dynamics:**
  How can people-side KM insights be translated into continuous learning systems that support adaptive, organization-wide AI adoption?

This study integrates knowledge management and semantic governance (information architecture, metadata, provenance), organizational network science (ONA and diffusion), socio-technical and organizational learning theories, behavioral insights for adoption, and emerging enterprise agentic-AI governance and design. It shows how behavioral and relational

signals can be translated into privacy-preserving, machine-usable indicators (e.g., trust density, connector density, knowledge latency) that make knowledge flow and adoption readiness observable at scale. These integrated perspectives motivate the framework developed in this study.

This study draws on prior scholarship, industry leader commentary, practitioner interviews, and large-scale behavioral and collaboration data generated through LEAD, a knowledge-sharing platform founded by the author. LEAD interaction data support the modeling of organizational networks and work patterns across more than 2,000 organizations. The organizations represented in this dataset typically operate at or above 500 employees or exhibit comparable structural complexity, such as multiple teams or business units. These characteristics align with enterprise environments discussed in organizational learning and knowledge-management literature (Dalkir, 2017; Nonaka & Takeuchi, 1995) and in industry classifications of enterprise-scale firms (Gartner, 2025). Although the specific tools used by these organizations vary and are not part of the study data, enterprises of this size generally rely on digital collaboration and knowledge systems that support coordinated work and cross-boundary information flow. As a result, the behavioral and network patterns observed here provide a relevant window into how knowledge moves in enterprise contexts, while recognizing that organizational context and sector-specific practices vary.

Synthesizing these findings, this study introduces the Organizational Intelligence Loop (OIL). Because enterprise GenAI builds and operates on an organization's existing knowledge base, its effectiveness depends on how well that knowledge is structured, connected, and machine readable. OIL builds on this premise, treating knowledge as a "living infrastructure" that links human behavior, information flow, and system feedback so AI can learn and scale effectively. The framework integrates three dimensions:

1. How people collaborate and influence one another,
2. how information remains transparent, governed, and trusted (including organizational context such as identities, access boundaries, and decision rights),
3. How behavioral nudges and reinforcement mechanisms embedded in daily work drive adaptation over time.

By operationalizing these dynamics, OIL positions knowledge management as an adaptive system for enterprise intelligence that supports both AI scaling and broader organizational redesign. In doing so, it reframes AI adoption from a one-time deployment challenge into an ongoing, measurable learning process.

To examine the existing evidence on these questions systematically, the next section reviews the behavioral, network, and knowledge-management foundations that shape enterprise AI adoption.

Literature Review & Relevant Research[1]

This section reviews the current state of the research on AI adoption, examining published articles and insights derived from a large LEAD dataset encompassing more than 2,000 organizations.

Part 1: Shifting the Debate from Access to Agency in Enterprise AI Adoption

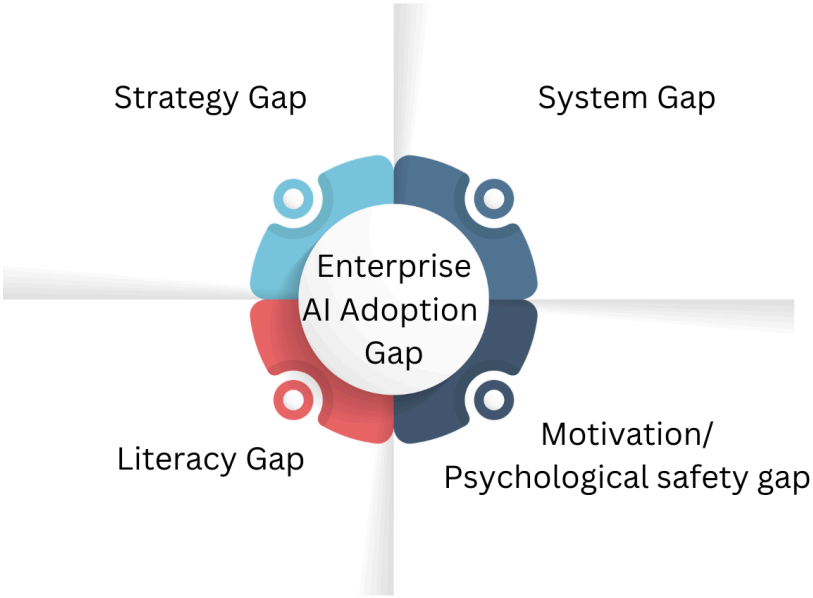

Despite rapid advances in generative and agentic AI, adoption across enterprises remains uneven with large differences between organizations in both speed and realized value, and often underwhelming in impact. (Menlo Ventures, 2024). Organizations with clear strategic intent and experimentation frameworks integrate AI faster and more productively than others which respond with restrictive governance, throttled access, or fragmented pilots (McKinsey & Company, 2025). Recent analyses suggest that firms lacking a defined AI rationale and success metrics underperform because they fail to design and test the social and knowledge systems that enable change (McKinsey & Company, 2025).

A related factor shaping uneven enterprise outcomes is technology maturity, encompassing both the maturity of AI systems and the organization's readiness to operationalize them. Early-stage GenAI tools frequently outpace the structures required to deploy them responsibly, creating gaps between theoretical capability and what the organization can actually support in practice (McKinsey & Company, 2024). This results in a situation where the "technology curve" accelerates faster than the governance, culture, process, and decision-logic can absorb the technology (McKinsey & Company, 2025).

[1] Article in prep

Recent KM and AI governance work also suggests that what must be governed is expanding beyond domain content and policies to include the human context through which knowledge is created, interpreted, and acted upon. Over the last decade, KM has increasingly emphasized that knowledge is context-dependent and embedded in human interaction rather than reducible to repositories alone (ISO, 2018). Similarly, the NIST AI Risk Management Framework highlights trust, context, and human dynamics as central to responsible AI adoption (NIST, 2023).

Socio-technical systems theory argues that organizational performance emerges from the joint optimization of human activity and technical infrastructure (Trist & Emery, 1951; Baxter & Sommerville, 2011). Distributed cognition similarly emphasizes that knowledge is distributed across people, tools, and representations rather than located in individuals alone (Hutchins, 1995). Network science extends this view by showing that connectivity, influence patterns, and structural bottlenecks shape whether knowledge and new practices propagate or stall (Freeman, 1978; Borgatti et al., 2018; Valente, 2012). Together, these perspectives imply that enterprise AI adoption requires governance not only of domain knowledge and policies, but also of the behavioral and network context that determines how knowledge is created, trusted, and acted upon.

At the same time, “AI literacy’, a commonly proposed remedy, remains underspecified and often insufficient. Most AI literacy programs teach employees how to prompt tools or interpret outputs, but they rarely address enterprise-specific knowledge flows, contextual judgment, tacit knowledge, or the psychological and behavioral conditions that determine whether AI recommendations are trusted or acted upon (Edmondson, 2019). Recent analyses from industry and academia similarly note that organizations often overemphasize technical upskilling, while neglecting the social and contextual expertise required to use AI effectively (McKinsey & Company, 2024, Orlikowski & Scott, 2014).

Research in human AI teaming shows that effective use hinges not only on skill but also on cognitive framing (Hutchins, 1995), situational awareness, and trust calibration in human–AI interaction (Kosmyna et al., 2025). This suggests that the barrier is not simply “literacy,” but deeper organizational insight gaps, foreshadowing the argument developed in the literature review Part 3.

Most scholars correctly identify these challenges as issues of governance and capability development, but tend to frame them in structural or policy terms (Brynjolfsson & McAfee, 2017). Far less attention has been paid to how human dynamics, including trust, influence, psychological safety, mediate the perceived legitimacy and ultimate effectiveness of AI tools (Edmondson, 2019). Research in distributed cognition further shows that heavy reliance on automation can inadvertently offload cognitive work, reducing situational awareness and degrading decision quality (Hutchins, 1995). This dynamic underlies warnings that overreliance on AI without robust cognitive safeguards can erode human reasoning and creativity, a mechanism consistent with emerging concerns about cognitive offloading–driven degradation and “distributed atrophy.”(Hutchins, 1995; Kosmyna et al., 2025).

These insights shift the debate from access to agency, highlighting that the challenge is not just whether employees can use AI, but how they retain autonomy, judgment, and learning capacity as AI becomes embedded in daily work. In other words, the central question becomes how organizations design conditions under which people can exercise informed agency with AI, rather than simply granting them tool access.

Part 2: Behavioral and Network Foundations of Technology Adoption

2.1 Behavioral Insights: Psychological Mechanisms Shaping Adoption

“Behavioral Insights” is grounded in the idea that human behavior is shaped by autonomic responses (Kahneman, 2011). These are rapid, nonconscious reactions to stimuli that occur before deliberate reasoning begins. They include instinctive approach or avoidance tendencies, quick judgments based on familiarity, and immediate impressions of whether information or a messenger is trustworthy (Kahneman, 2011). Cognitive biases and social norms, which are part of this construct, further guide decision-making by shaping what individuals notice, how they interpret information, and which actions feel socially acceptable (Cialdini, 2007; Thaler & Sunstein, 2008).

Nudge Theory, introduced by Thaler and Sunstein (2008), builds on these observations by showing how small contextual adjustments can influence behavior without restricting choice. Nudges shape behavior because people rely heavily on “mental shortcuts,” emotional cues, and group expectations when navigating uncertainty. In organizations, nudges can shift how employees engage with new technologies by influencing what feels easy, normal, or socially supported.

A complementary framework widely used to categorize these nonconscious influences is MINDSPACE, which refers to Messenger, Incentives, Norms, Defaults, Salience, Priming, Affect, Commitment, and Ego. MINDSPACE brings together behavioral insights in an easy to remember acronym to help business leaders explain how subtle contextual features shift behavior in systematic ways (Hallsworth & Kirkman, 2020). Messenger credibility, default settings, the visibility of norms, emotional tone, and the perceived relevance of information all shape whether employees view emerging technologies as safe, meaningful, and aligned with their group’s values.

Understanding these mechanisms for our study requires clarifying one way that norms play out: “social proof.” Social proof refers to the tendency for individuals to mimic the actions of others (to follow norms) when determining what is appropriate or trustworthy, particularly under uncertainty (Cialdini, 2007). In organizations, visible peers experimenting with a new technology create evidence that the behavior is safe, expected, and valued. This visibility contributes to the formation of new “norms,”meaning shared expectations about what behaviors are typical (descriptive norms) or desirable (injunctive norms) within a group (Cialdini, 2007; Cialdini & Goldstein, 2004*).*

While these behavioral mechanisms explain individual-level tendencies, an important extension concerns how they scale through influence and social structure. Influence is unevenly distributed within organizations, and certain individuals hold positions that magnify the reach of their attitudes and behaviors. Here, research on network interventions becomes essential. Valente’s network intervention framework shows that relationship patterns can be intentionally shaped to accelerate behavioral diffusion(Rogers, 2003; Valente, 2012*)*. For example, leaders

can identify central influencers, strengthen ties to isolated groups, or create bridging relationships that allow new practices to spread more rapidly. These interventions demonstrate that network structure is not fixed; this malleability signals an organization's action potential for change and can be nudged to improve the likelihood of successful adoption (Valente, 2012).

Within knowledge management and organizational network analysis (ONA), those "central influencers" are described as connectors, brokers, or boundary spanners. They link separate groups, transfer tacit knowledge, and shape collective interpretation of facts (Tushman & Scanlan, 1981; Burt, 1992; Cross & Parker, 2004). These individuals function as behavioral conduits through which MINDSPACE and Nudge Theory mechanisms travel. Their credibility as messengers, their position within visible norms, and the relational pathways that support commitment and recognition are likely to influence whether innovations such as AI gain traction or stall.(Dolan et al., 2012)

Understanding behavioral insights at both the psychological level and the network-structural level clarifies why influence interventions are a powerful driver of technology adoption. New behaviors spread not only because information exists, but because credible and well-positioned individuals model those behaviors, reinforce them, and occupy structural positions that allow their actions to ripple throughout the organization.

2.2 Network Structures and Diffusion: How Influence Scales

Building on these behavioral foundations, extensive network research demonstrates that invisible social structures, remote ties, informal advice channels, and influence patterns strongly predict how knowledge, innovation, and technology adoption spread within organizations (Cross & Parker, 2004; Burt, 1992). Connectors are highly central actors with many direct ties, while brokers or boundary spanners link otherwise disconnected groups and bridge structural holes. In many cases the same individual performs both functions (Burt, 1992; Cross & Parker, 2004; Tushman & Scanlan, 1981). Connectors amplify diffusion within groups, whereas brokers enable ideas and practices to move across group boundaries. Conversely, bottlenecks, fragmentation, or cliquish networks can slow the spread of information and impede adoption (Freeman, 1978). Advances in Organizational Network Analysis (ONA) now allow these dynamics to be measured quantitatively through metrics such as betweenness centrality, collaboration density, reciprocity, cross functional connectivity, and knowledge latency (Freeman, 1978; Borgatti et al., 2018). These measurable indicators reveal the hidden architecture of collaboration and enable organizations to observe whether knowledge is spread, socially adapted, or at risk of stagnation.

It is noteworthy that not all network signals carry the same interpretive weight, and ONA research type distinguishes between behavioral trace data and perceptual or relational data. "Passive ONA" research captures digitally-observable collaboration patterns such as meeting frequency, communication traces, or workflow interactions, while "Active ONA" captures subjective relational qualities such as trust, reliance, perceived expertise, and influence as reported by individuals in the network (Cross & Parker, 2004; Borgatti et al., 2018). These two modes do not translate equivalently: passive signals reveal structural opportunity, whereas active signals reveal motivational and relational context (Cross & Parker, 2004; Borgatti, Everett, & Johnson, 2018). Together they offer a multidimensional view of diffusion potential, bottlenecks, and tacit knowledge pathways, extending earlier work by treating structural

opportunity and relational motivation as complementary lenses on how networks shape diffusion and adoption (Cross & Parker, 2004; Borgatti et al., 2018).

Network metric benchmarks further help define what healthy systems look like. Research in network science shows that diffusion depends primarily on the structure of the network itself—whether it contains sufficient density, bridging ties, and cross-boundary pathways to support propagation (Watts, 2011; Watts & Dodds, 2007). Within these structures, a subset of individuals often occupy positions that carry disproportionate collaboration load. Cross, Rebele, and Grant (2016) find that in many organizations, 3–5 percent of employees account for 20–35 percent of high-value collaborations, functioning as informal connectors who facilitate learning, decision-making, and change adoption (Cross, Rebele, & Grant, 2016; Cross & Parker, 2004). These individuals are not necessarily formal leaders; many sit in mid-level or individual-contributor roles. Their influence arises not solely from personal attributes but from their structural position in the network.

Critically, these actors do not "cause" diffusion independently. Their ability to amplify or transmit new practices depends on whether the surrounding network provides enough pathways for ideas to cascade. A small number of central actors may initiate or accelerate spread, but broad adoption requires distributed bridges and overall connectivity. For this reason, connector distribution becomes an important benchmark: when only a very small fraction of employees serve as cross-team bridges, innovation slows and teams become increasingly isolated (Cross & Thomas, 2011). Overreliance on a small number of well-positioned individuals can also create operational risk. When too many colleagues seek support from the same connector, overload and burnout increase, making connector saturation an early indicator of performance risk and potential turnover (Cross et al., 2016).

Complementary research in change management translates these network principles into practical frameworks. The eight-step model emphasizes coalition-building and visible wins (Kotter, 1996); "Nimble Networks" highlight the value of cross-unit brokers for accelerating collaboration (Hansen, 2009); and the ADKAR framework operationalizes individual readiness as a diffusion process (Prosci, 2006). Together, these perspectives reinforce that change spreads through social diffusion—the relational propagation of ideas and behaviors through observation, imitation, reinforcement, and trust—rather than through top-down instruction. Large shifts often emerge from small cascades once early influencers activate the right thresholds (Watts, 2011).

These network and behavioral principles are not unique to internal collaboration, as fact, they already underpin many business functions. In marketing and sales, for instance, *network effects* determine how products and ideas spread across customer ecosystems, where some customers hold disproportionately higher influence and propagation value than others (Andreessen Horowitz, 2023). Metrics such as lead-response time, follow-up cadence, and message re-sharing then measure how effectively organizations activate these influence systems and sustain engagement. These are, at their core, behavioral diffusion indicators: faster feedback loops and visible social proof increase trust and participation.

The same logic applies within organizations. Functions such as HR, Operations, and Engineering can adapt these behavioral metrics to track how quickly employees respond to new initiatives, share learnings, or reinforce adoption. By extending these "external" behavioral

playbooks inward, enterprises can create a unified system of measurable engagement that links customer, employee, and AI interaction dynamics (Watts, 2011).

Recent applications extend this understanding toward what some describe as *behavioral intelligence*. Collaboration signals such as meeting cadence, responsiveness, or communication patterns can correlate with overload, disengagement, or resistance (Cross & Parker, 2004; Dalkir, 2017). These data-driven methods detect emerging risks earlier than traditional surveys. However, most organizations still treat informal networks as secondary to formal hierarchy, which means that the individuals who carry disproportionate influence receive little recognition or support for the invisible work they perform. *(Cross, Borgatti, & Parker, 2002;)* This misalignment delays the identification of resistance, increases fatigue among key contributors, and weakens the informal exchanges through which most tacit knowledge is shared.

The AI adoption context may amplify these network dynamics in both positive and negative directions. In principle, well-connected environments with visible expertise and trusted influencers can accelerate diffusion, as employees model one another's experimentation and rapidly share emerging practices. However, most organizations do not operate with this level of network visibility or cohesion. In practice, limited awareness of expertise, influence, and trust relationships often leads to fragmented literacy-building efforts, uneven experimentation, and the expansion of '"shadow AI"practices (Gartner, 2025). Shadow AI refers to employees' unsanctioned use of external or non-compliant tools to compensate for restricted access or unclear governance. Although often well-intentioned, these practices bypass organizational safeguards, create security and compliance risks, and reinforce isolated pockets of experimentation. The result is a widening divergence between formal AI initiatives and actual day-to-day behavior, undermining data quality, governance, and trust in enterprise AI systems.

Although prior research acknowledges that social diffusion and trust shape technology uptake (Rogers, 2003; Cross, Parker, Prusak, & Borgatti, 2001), few studies directly link these behavioral and network dynamics to quantifiable knowledge-system metrics or to observable performance outcomes of agentic AI systems.

Part 3 : Defining People-Side Knowledge Management Infrastructure and Its Relationship to Other KM Layers

3.1 From Content-Centric KM to People-Side KM Infrastructure

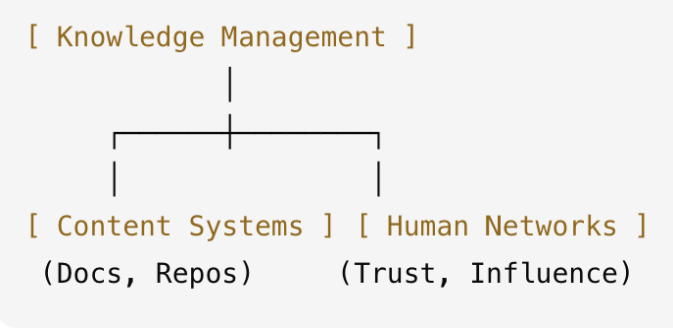

In most organizations, knowledge management (KM) still evokes images of repositories, intranets, and content systems. This view is rooted in codification and reuse (Davenport & Prusak, 1998; Dalkir, 2017) and remains essential for organizing explicit knowledge but increasingly insufficient for the adaptive learning required in AI-driven environments, even when strengthened by semantic and metadata practices (e.g., taxonomies/ontologies, normalization, and provenance standards) that improve retrieval and transparency but do not capture trust, influence, or tacit exchange pathways. *Adaptive learning* originally referred to how individuals adjust strategies in response to feedback (Corbett, Koedinger, & Anderson, 1997). At the organizational level, it describes how groups modify shared knowledge and behavior based on continuous input from the business environment, people and technology.

Foundational KM research established that productive knowledge systems rely not only on formal content but also on the social and relational networks through which knowledge flows (Nonaka & Takeuchi, 1995; Wenger, 1998). Building on this principle, organizations adopting AI at scale must treat these human networks as part of their core knowledge infrastructure rather than as informal by-products of collaboration (Cross & Parker, 2004; Wenger, 1998; Nonaka & Takeuchi, 1995). *People-side KM infrastructure* refers to the systems, data, and processes that capture and strengthen these social dynamics. However, this capability presupposes a baseline architectural substrate in the underlying information environment: coherent integration across collaboration, workflow, and knowledge systems; stable identity and access controls; and reliable data pipelines that preserve metadata and provenance (Ross, Weill, & Robertson, 2006; Weill & Ross, 2004; Khatri & Brown, 2010). Where these conditions are absent, organizations face fragmentation, inconsistent signals, and reduced confidence in analytics- or AI-supported outputs—making information-architecture modernization a prerequisite rather than an AI-adoption “intervention.” (Wang & Strong, 1996; Strong, Lee, & Wang, 1997). Examples include trust ties, collaboration patterns, and influence relationships. Such systems, data and processes must be measured, governed, and improved alongside other knowledge assets.

3.2 Operationalizing People-Side KM: Active and Passive ONA

The literature review Part 2.2 introduced the distinction between passive ONA, based on behavioral trace data, and active ONA, based on perceptual and relational data, emphasizing that the two types do not translate directly and cannot be interpreted interchangeably (Cross & Parker, 2004; Borgatti, Everett, & Johnson, 2018). This distinction becomes even more

important when ONA is treated not merely as an analytical technique but as a foundational component of knowledge infrastructure.

However, there is a gap between research evidence on the distinctive benefits of active, perceptual network data and practitioners' tendency to rely primarily on passive digital traces. Many academic KM studies rely almost exclusively on active ONA, reflecting the field's qualitative origins in advice seeking, trust, and perceived expertise (Cross & Parker, 2004). By contrast, engineers and data practitioners, such as those publishing from large platforms and enterprise tooling teams, tend to rely on passive ONA because it scales automatically across digital work environments. A third group blends the two, often without clarifying how the signals differ or how the underlying assumptions shift when moving from self-reported ties to behavioral edges (Contractor, Wasserman, & Faust, 2006; Borgatti, Everett, & Johnson, 2018)..

This gap produces several challenges. First, the volume and density of edges differ dramatically between passive and active networks: behavioral data may generate thousands of weak or incidental edges, while relational data may reveal only a few strong ties (Borgatti et al., 2018). Without clear interpretive guidance, these networks appear incompatible or contradictory, even though they reflect different layers of the same system (Cross & Parker, 2004; Borgatti et al., 2018).

Second, terminologies such as node, edge or tie, attribute, ego network, and directed versus undirected or open networks vary widely across fields and authors (Borgatti et al., 2018). As synthesized by (Pugh, 2023), KM practitioners use less social graph terminology differently than organizational network researchers, adding further variation to how concepts such as ties, directionality, weight, and role attributes are defined. This inconsistency complicates attempts to standardize knowledge-flow metrics or compare findings across studies.

Third, because many organizations analyze only one type of ONA, they lack a coherent way to connect behavioral opportunity, "who interacts?", with relational motivation, "Who is trusted or relied upon?". Treating these signals as interchangeable produces misleading interpretations, such as assuming that high communication frequency implies high trust or equating workflow dependency with expertise attribution (Cross & Parker, 2004). These assumptions rarely take into account that frequent communication may arise from coordination demands, group meetings to which someone is merely invited, large distribution-list emails, or informal social chatter rather than genuine reliance. Likewise, the structural dependency data in passive ONA do not necessarily mean that an individual is perceived as credible or expert.

For the people-side KM infrastructure, this distinction is not merely theoretical. It determines how organizations measure tacit knowledge flow, identify credible messengers, map expertise pathways, anticipate overload, and design AI-enabled knowledge systems (Cross & Parker, 2004; Borgatti et al., 2018). Clarifying how active and passive ONA differ, and how they must be integrated, provides a conceptual foundation for a unified relational layer that can support adaptive learning, change readiness, and AI adoption at scale.

3.3 ONA as a Mechanism Within Knowledge Systems

Building on the foundations established in Parts 2.2 through 3.2 and consistent with patterns observed across contemporary KM and network research, ONA can be understood not only as

a diagnostic technique but also as a mechanism that makes the underlying architecture of knowledge movement measurable and governable within organizational systems (Cross & Parker, 2004;). Within this synthesized perspective, ONA functions as more than an analytical tool. It operates as an dynamic infrastructural layer that links human interaction patterns with the performance of knowledge systems at scale (Cross & Parker, 2004; Monge & Contractor, 2003). This framing reflects a broader trend in the literature in which relational dynamics, behavioral context, and structural connectivity are increasingly treated as core components of how organizations interpret, manage, and leverage knowledge flows(Cross & Parker, 2004; Davenport & Prusak, 1998).

Immature enterprise knowledge systems have been built around document-centric intelligence—repositories, intranets, dashboards, and, more recently, LLM-indexed content—emphasizing codified information rather than the human context required to interpret it (Nonaka & Takeuchi, 1995; *Dalkir, 2011, Milton & Lambe, 2019*). KM research has long shown that documentation trails and explicit knowledge artifacts capture only a portion of organizational knowledge, omitting tacit expertise, practical judgment, and the social pathways through which knowledge actually moves (Nonaka & Takeuchi, 1995; *Baumard, 1999; Stepper, 2015*). Consequently, many enterprise AI and internal search systems underperform not because documents are inaccurate, but because the organizational context that gives them meaning is missing (Kellogg et al., 2020)

ONA provides this missing context by operationalizing several categories of insight consistent with established network research (Cross & Parker, 2004; Borgatti, Everett, & Johnson, 2018).

Knowledge captured through digital exhaust in Passive ONA
Passive ONA quantifies how work moves by measuring structural patterns such as density, distance, clustering coefficient, and effective size. These indicators reveal whether collaboration occurs in closed networks with high redundancy, open networks that enable cross-team learning, or structures where one-node structural holes create bottlenecks in knowledge movement. Metrics such as in-degree, out-degree, and betweenness centrality highlight where information flows freely versus where it becomes delayed or distorted across boundaries (Cross & Parker, 2004; Borgatti et al., 2018).

Relational and motivational context in Active ONA
Active ONA is generated by a survey administered to individuals in a network, who, in turn, assess their network members and connections. In addition to the metrics calculated with Passive ONA, active ONA surfaces interpersonal dynamics that shape whether information will be accepted or acted upon. Measures such as reciprocity, tie strength, homophily, and propinquity, while also available in Passive ONA, are supplied by participants, and thus provide visibility into trust flows, perceived credibility, and social cohesion. These relational attributes, coupled with survey respondents' judgment and values, determine whether a message spreads, stalls, or is ignored—insight that cannot be derived from documents or search logs alone.

Tacit knowledge pathways.
Integrating passive and active ONA helps identify brokers, bridges, informal experts, and other high-influence individuals who hold essential tacit knowledge, but remain invisible in formal systems (Cross & Parker, 2004). Mapping their ego networks and cross-boundary ties supports expertise routing, onboarding, and knowledge-risk mitigation. Here, ego networks refer to the network formed by a focal individual ("ego"), all actors directly connected to that individual

("alters"), and the ties among those alters (Wasserman & Faust, 1994; Borgatti, Everett, & Johnson, 2018). These local networks provide a lens on the informal influence, support, and information flows that surround individuals in organizations (Borgatti et al., 2018).

Cross & Parker (2004) point to scenarios where an individual's departure (e.g., through re-assignment, layoff, attrition, etc.) results in operational holes (Gamba et al., 2024; Younis, 2023; Phaladi, 2022). For example, at a large global consumer goods manufacturer, the webmaster's reporting routines incorporated interpersonal ties and project awareness between the webmaster and her clients. When she moved, web usage dropped as promotion was not connected to the operational reality.

The analysis might also show that some tacit knowledge has no effective pathway. In many organizations, expertise remains trapped in isolated individuals or small clusters with few reciprocal ties or limited cross-functional exposure. Valente's diffusion research demonstrates that when relational channels are absent or fall below the thresholds needed for diffusion, knowledge fails to propagate even when it is valuable (Valente, 2010; Watts & Strogatz, 1998). Identifying both existing and missing pathways is therefore essential for understanding where tacit knowledge can move and where it cannot.

Behavioral knowledge acquisition.

As enterprise AI moves from retrieval to reasoning and decision support, it requires visibility into how work actually happens, not merely what is documented. This paper uses the term *behavioral knowledge acquisition* to describe the continuous capture and interpretation of signals that reflect the lived dynamics of collaboration. Examples include:

- Interaction patterns that reveal the rhythm and consistency of knowledge exchange
- Indicators of expertise such as repeated reliance from specific groups or functions
- Trust-confirmed relationships that shape whether information will be accepted
- Influence positions, understood as network positions whose effect depends on the surrounding structure and whether subsequent layers of the network will propagate new practices
- Knowledge-transfer events that identify where tacit information moves or stops
- Cross-functional dependencies that show how workflows propagate across boundaries
- Continuity risks such as when critical knowledge holders leave, transfer internally, or approach retirement without a clear succession pathway
- Group coordination and behavioral teaming dynamics that reveal how teams actually form and collaborate in practice that reveal how groups coordinate in practice, beyond formal skills, including responsiveness, willingness to collaborate, informal leadership, reciprocity, and psychological safety

The outsized influence of these signals suggests that skills alone do not determine whether a team can execute effectively. Skill gaps describe what a group lacks in principle. Behavioral gaps describe what it cannot accomplish in practice. AI systems and human managers often overlook these dynamics, leading to teams that appear complete on paper but struggle in real workflows (Cummings & Pletcher, 2011).

Interpreting these signals made available through ONA aligns with socio-technical KM perspectives that view behavior, relationships, and context as essential dimensions of

organizational knowledge (Nonaka and Takeuchi, 1995; Davenport and Prusak, 1998). These patterns make visible where tacit knowledge resides today and where it may disappear tomorrow.

Within this framing, ONA becomes an analytical engine that translates structural and relational attributes into machine-usable representations of organizational reality (Borgatti et al., 2018). AI systems trained only on textual or structured content lack this grounding. They frequently misidentify experts, overlook hidden dependencies, or propose actions that contradict lived collaboration patterns (Wang, 2015). These findings provide empirical support for treating people-side KM as core infrastructure and motivate the implications and research questions developed in Part 5.

AI Agent Workflows at Risk

The stakes increase when organizations deploy AI agents or digital twin workflows. These systems attempt to emulate human decision processes, yet their effectiveness depends on accurate representations of who participates in which decisions, where bottlenecks form, and how coordination unfolds across teams. When behavioral knowledge is missing, digital twins may omit key actors, misrepresent succession risk, or fail to detect that a critical knowledge holder is about to transfer or depart. Such blind spots often remain invisible until after disruptions occur, making behavioral knowledge essential for resilient automation and responsible AI design.

Applications of ONA further demonstrate the value of behavioral knowledge. ONA does not prescribe strategic decisions; it provides empirical evidence about how work and knowledge actually move through the organization. In practice, ONA can surface hidden experts relevant to onboarding, quantify structural holes and dependency risk that strengthen the case for targeted knowledge transfer, identify influential nodes and trust hubs that inform change-readiness planning, and document misalignment between formal structures and observed collaboration patterns—evidence leaders can then interpret alongside business objectives, constraints, and governance requirements (Cross & Parker, 2004; Kotter, 1996). These insights cannot be derived from documents alone; they require the integration of behavioral signals with structural and relational attributes.

As established in Part 1, leading KM and AI governance frameworks (ISO 30401; NIST AI RMF) treat context, trust, and human dynamics as core requirements for responsible deployment. Here, behavioral knowledge acquisition operationalizes those requirements by supplying continuous, organization-specific signals about trust, influence, expertise, and continuity risk—inputs that agentic workflows typically lack.

These foundations imply that network and behavioral evidence matter for two reasons: they shape implementation trajectories and they determine whether agentic systems reflect real coordination pathways. The next section turns to practitioner-generated evidence from large-scale enterprise systems that instrument these dynamics directly.

Part 4: Practitioner Evidence From Large-Scale Enterprise ONA Systems (LEAD Case Studies)

In parallel with academic work, practitioner platforms that instrument real-time enterprise collaboration have begun generating empirical evidence that enriches KM and ONA research. These large-scale systems capture behavioral and structural patterns that are rarely observable in traditional studies, providing complementary insight into how knowledge actually flows within modern organizations.

The LEAD platform is one such example. LEAD was developed by the author's organization. It offers a large , multi-year dataset on enterprise behavioral knowledge flow. Its large-scale observations reinforce and extend the literature reviewed earlier on knowledge networks, social capital, and people-side KM.

4.1 Overview of LEAD as a Practitioner Research Source

LEAD is an enterprise knowledge-management and organizational-network system used by more than 2,000 organizations globally. The platform integrates passive ONA data inputs, such as collaboration metadata from Slack, Microsoft Teams, Google Calendar, and Outlook, with active ONA inputs, including targeted surveys, employee-matching interactions, engagement patterns, and workflow telemetry. Sample metadata includes calendar co-attendance and recurring meeting patterns, Slack/Teams interaction structure (e.g., channel co-membership), online/offline activity trends, and—when provided by the organization—basic node attributes such as business unit, geography, role level, and tenure. These fields allow interactions to be represented as a graph of nodes (employees) and edges (observed collaboration ties) for organizational network analysis (ONA), a workplace application of social network analysis (SNA).

Across these deployments, LEAD applies graph-theoretic methods and network modeling, operationalized through ONA alongside behavioral analytics, to map how knowledge actually moves inside organizations, including who knows what, how information flows, and where social or structural barriers emerge. Although many organizations adopt LEAD for people-focused programs such as coworker introduction, buddy systems, or onboarding, these interaction data also generate evidence of trust ties, contextual learning, and informal expertise exchanges.

These signals map to common ONA/SNA measures used in prior research, including both network-level structure and individual (node-level) centrality:

- **Network Density**: proportion of realized ties, indicating overall connectedness.
- **Tie strength**: frequency or intensity of dyadic interactions. This is relevant to tacit knowledge exchange.
- **Reciprocity**: mutuality of ties, often linked to trust and psychological safety.
- **Clustering coefficient**: the degree to which a person's contacts are connected to each other, shaping redundancy and localized cohesion.

- **Average Betweenness centrality**: the extent to which actors occupy bridging roles across clusters, enabling access to nonredundant knowledge (Pugh, 2023).
- **Proximity to executives:** an individual's network distance to senior leaders (e.g., shortest-path steps), used as a proxy for access to decision-makers and influence pathways.
- **Proximity to connectors:** an individual's network distance to high-centrality "connectors" (e.g., top-percentile degree/betweenness), indicating how easily they can reach key brokers for information flow and coordination.

These metrics show where knowledge flows smoothly or gets blocked, which people bridge structural gaps, and where informal trust relationships foster or hinder learning.

4.2 Dataset Scope

LEAD has been deployed in 2,000+ organizations. This study analyzes a 2020–2025 subset of approximately 1,500 organizations for which longitudinal time coverage and analysis permissions are consistent. Within this analytic subset, additional cohorts are defined by research design. Case Study 1 draws on a panel of approximately 800 organizations (2020–2022) to examine early-stage adoption patterns and attrition among informal influencers. Case Study 2 situates two focal organizations within the broader ~1,500-organization (2020–2025) cross-section, allowing their network maturity and outcomes to be interpreted against a wider benchmark.

Across these deployments, the dataset includes:

- 4M calendar events (meeting patterns and participation)
- 300M presence indicators (responsiveness and availability)
- 1M peer-connection and mentorship links
- 60K survey responses (sentiment, perceived influence, trust)
- 4M collaboration-channel memberships (cross-functional overlap)
- 160K anonymized role/tenure/centrality profiles
- 135K behavioral-tone messages (metadata only; no content)

Additional data sources include product-usage metrics, system logs, and other behavioral attributes not fully represented in the list above, such as broader patterns of interaction frequency, working rhythms, and role-based network position. Consistent with responsible KM and analytics practice, personally identifiable information (PII) was excluded; all data were aggregated or anonymized under SOC 2–aligned governance standards, ensuring that individual-level identifiers were neither stored nor analyzed.

4.3 Derived Constructs and the Integration of Passive and Active ONA

When combined, the structural and relational indicators captured through LEAD—using established network metrics above —can be interpreted as forming a behavioral knowledge graph that models how knowledge circulates within organizations (Borgatti, Everett, & Johnson, 2018; Burt, 1992; Cross & Parker, 2004).

As detailed in Parts 2.2 and 3.2, two practical gaps limit how prior KM and ONA scholarship translates into modern enterprise systems:

Translation gap: There is no commonly accepted framework for translating passive ONA traces (e.g., interaction patterns, temporal metadata) into constructs that historically rely on active ONA, including trust, perceived expertise, and willingness to seek advice.

Scale-and-timeliness gap: Earlier studies were constrained by the data available at the time and did not fully address the analytical challenges introduced by continuous, large-scale behavioral traces in contemporary enterprise platforms.

To address these gaps, the author's company, LEAD, has introduced several derived constructs that extend classical ONA metrics with relational or contextual qualifiers relevant to KM. These constructs do not replace standard measures but clarify aspects of collaboration that structural or survey-based ONA alone cannot capture (Burt, 1992). Collectively, these extensions allow the underlying behavioral, structural, and temporal inputs to be represented as a more multidimensional behavioral knowledge graph, one that links relational reliability, information velocity, and collaboration health—all of which shape organizational intelligence.

Nonresponse-Aware Integration of Passive and Active ONA

In addition to bridging passive and active ONA signals, LEAD operationalized this nonresponse-aware approach independently as a production necessity, and it is consistent with prior survey-methods research showing that response-process metadata can classify meaningful nonresponse patterns and inform inference under uneven participation *(Bosnjak & Tuten, 2003; Biemer et al., 2013).*

- **Nonresponse is treated as a behavioral signal, not unusable missing data.**
- **Nonresponse is differentiated by pattern** (e.g., one-time nonresponse, persistent nonresponse, "viewed/opened multiple times but did not respond").
- **Nonresponse is interpreted in context using auxiliary signals** (e.g., role level, tenure, network position/centrality, workload proxies, response history, observed collaboration intensity).
- **Confidence varies with evidence of strength.** Where context is limited, LEAD treats interpretations as lower-confidence; where converging signals exist, inferred labels can be treated as stronger evidence for downstream constructs.

This makes derived constructs usable at scale. By distinguishing nonresponse types and leveraging cross-signal context, LEAD can compute relationship- and trust-adjacent indicators more consistently in real enterprise environments where survey completion is uneven and systematically biased.

Example construct: Trust density

One example used in this study is trust density, defined as:

Trust Density = Number of trust-confirmed ties / Total number of possible or observed ties

- The numerator draws from active ONA, reflecting ties where individuals explicitly confirm trust or reliance (e.g., survey-confirmed trust relationships).

- The denominator draws from passive ONA, representing the set of observed collaboration ties (e.g., repeated calendar co-attendance, recurring meeting patterns, Slack/Teams structural interaction such as channel co-membership) that indicate opportunity for knowledge flow but do not confirm whether trust-based exchange actually occurs.

Trust density therefore distinguishes between **structural opportunity and relational reliability,** aligning with established findings that trust shapes whether information is accepted, reinforced, or ignored (Davenport and Prusak, 1998; Nonaka and Takeuchi, 1995), and preserves the conceptual separation between structural opportunity and relational reliability. As a result, this complements "clustering," "density," "reciprocity," and "distance" metrics recognized in KM and ONA scholarship (Borgatti et al., 2018; Cross and Parker, 2004).

In practice, **trust itself is not monolithic:** Relational reliability varies by domain, meaning individuals may be trusted in one area of expertise but not in another. Social capital research shows that people differentiate between task-based trust, knowledge-based trust —- either experience-based or reputation-based —-, and affective trust, each of which governs different forms of knowledge exchange. This distinction matters because a structural tie may support high trust for technical problem solving while offering little reliability for guidance in unrelated domains. As a result, organizations often exhibit multiple "trust densities," each linked to a specific knowledge domain or context of practice. Recognizing this variation preserves the conceptual separation between structural opportunity and domain-relevant relational trust, and aligns with findings that knowledge flows only when trust is appropriate to the content being exchanged. Additional conceptual illustration of measurable indicators derived from integrated ONA appears in A3 within Appendix A.

Other derived constructs follow this logic. Node centrality measures (e.g., in- or out-degree, betweenness, eigenvector) may be qualified by perceived expertise, and cross-boundary ties can be interpreted in terms of accessibility or observed advice-seeking behavior. Because classical ONA terminology was not designed for systems that integrate behavioral traces with active relational signals at scale, this study introduces these derived constructs where needed. Each construct is defined when it first appears and computed within each organization's network. In LEAD's implementation, high-value nodes are treated as context- and time-dependent: individuals may become more or less central as projects shift, domains

change, or trust relevance varies over time. This operational assumption is also consistent with empirical findings that workplace ties are created, dissolved, and reactivated in response to changing conditions (Wu et al., 2021; Wu et al., 2023).

Together, these extensions provide a more complete view of collaboration health and people-side knowledge flow, complementing constructs recognized in KM and ONA scholarship (Borgatti et al., 2018; Cross and Parker, 2004).

4.4 Findings from LEAD Case Studies

Evidence from LEAD's large-scale deployments provides empirical support for themes emphasized in KM and ONA scholarship: informal relationships, trust dynamics, and everyday behavioral patterns significantly shape how knowledge moves within organizations. The cases summarized in Appendix A, A1 through A3 illustrate how these dynamics influence learning, continuity, and organizational readiness for new tools and processes.

Case Study 1: Informal Influencers and Knowledge Continuity (Appendix A, A1)

Across approximately 800 organizations observed between 2020 and 2022, individuals who voluntarily initiated peer-connection or onboarding activities frequently acted as informal brokers who linked otherwise disconnected groups. Their role resembled the bridging positions described in social-capital and diffusion research (Burt, 1992; Cross & Parker, 2004; Valente, 2010). When these contributors were not supported through sponsorship, resources, or integration with formal workflows, the informal programs they created often deteriorated.

Within 90 days, approximately 80% of these early contributors showed corporate email deactivation (bounce signals), and 48.99% had LinkedIn-verified job transitions; taken together, these indicators suggest substantial early-stage attrition in this cohort. Their departure coincided with the disappearance of the relational pathways through which these contributors had been transferring tacit know-how across employees differentiated by department, location, seniority, and role context, including interactions the organization would not typically capture in formal systems. This pattern mirrors long-standing KM concerns about the fragility of tacit knowledge when it depends on a small number of informal actors (Davenport & Prusak, 1998; Nonaka & Takeuchi, 1995) and highlights retention of informal influencers as a structural determinant of KM and AI readiness rather than a narrow HR outcome.

Case Study 2: Network Maturity, Psychological Safety, and Activation Speed (Appendix A, A2)

Appendix A, A2 contrasts two organizations that adopted LEAD.bot between 2021 and 2024. The first organization began with fragmented informal connectivity, low connector density, and limited peer learning. Improvements in responsiveness, cross-team collaboration, and execution emerged only after informal networks were intentionally strengthened and embedded into daily workflows through behavioral nudges such as lunch-and-learns, internal storytelling, and visible recognition for cross-team introductions.

The second organization entered with strong baseline connectivity, a consulting-oriented operating model, and high psychological safety. Employees were already accustomed to assembling new teams, working across functions, and integrating quickly into client environments. As a result, the program scaled almost immediately with near-universal participation and rapid response cycles, even without extensive internal promotion.

A notable contrast emerged during the transition from pilot to full deployment. Both organizations ran six-month pilots before annual subscription. Company A required almost three months of internal campaigning before cross team connections formed at scale, whereas Company B achieved widespread participation in the first week. This difference reflects both the pace of adoption and the pre existing level of reciprocity that became observable once the tool was introduced. As Edmondson (1999) suggests, psychologically safe environments treat new coordination practices as normal, while lower safety settings require sustained socialization and endorsement before discretionary collaboration emerges.

Psychological safety further shaped how anomalies were interpreted. The same minor timing issue caused by a daylight-savings shift led to urgent escalation and fear of sanction in Company A but was treated as inconsequential in Company B. This contrast is consistent with research showing that psychological safety influences whether deviations are seen as threats or as learning opportunities (Edmondson, 1999; Schein, 2010). Together, these observations reinforce a central conclusion in the network literature: the maturity of informal networks and the level of psychological safety significantly influence how effectively organizations learn, coordinate, and adopt new practices (Burt, 1992; Cross & Parker, 2004).

Across these cases, tool adoption did not unfold linearly but followed activation-curve dynamics. Collaboration remained limited until teams crossed a threshold of informal network maturity and psychological safety, after which engagement accelerated rapidly. Organizations with stronger baseline connectivity activated within days, whereas more fragmented environments required prolonged socialization before discretionary collaboration emerged. These patterns suggest that relational readiness, rather than mere exposure to the tool, largely shapes the pace at which new practices take hold, while recognizing that the evidence is observational and may reflect situational factors specific to each deployment.

Behavioral Regularities and Measurable Indicators (Appendix A, A3)

Appendix A, A3 summarizes insights from more than 1,500 organizations observed between 2020 and 2025. While it omits proprietary technical detail, it shows at a conceptual level how passive ONA, active ONA, temporal rhythms, and interaction patterns can be translated into measurable indicators of organizational health. The intent is not to disclose underlying algorithms but to demonstrate that such translation is feasible and analytically useful.

First, the aggregated evidence confirms that organizational performance, retention, and AI readiness are shaped less by formal structure and more by measurable behavioral dynamics.

- Collaboration patterns, trust-confirmed ties, and responsiveness cycles can be quantified and used as indicators of knowledge flow and coordination quality.
- Constructs such as trust density and knowledge latency, discussed earlier in the literature review Part 4.3 and elaborated conceptually in Appendix A, A3, illustrate how structural opportunity (for example, who appears together in meetings or shared

channels) can be distinguished from relational reliability (for example, who is actually trusted or relied upon), responding to a long-noted gap in KM and ONA literature regarding the operationalization of rich behavioral signals.

Second, the appendix highlights team-level regularities and social constraints.

- Knowledge-latency patterns, response speeds, and execution rhythms tend to converge within teams or functions, even when individual motivation varies, echoing prior research on shared temporal norms and group-level performance patterns (Waller, Zellmer-Bruhn, & Giambatista, 2002; Hackman, 2002; Edmondson, 1999).
- LEAD's observations show that a small subset of cultural champions or bridging contributors may operate faster than peers, but their behavior is still shaped by local norms and structural context (Cross & Parker, 2004; Burt, 1992).
- At the same time,case-level evidence suggests that when organizations identify such influencers and reinforce their behaviors through nudges and recognition, these individuals can gradually shift team-level dynamics, aligning with research on proactive socialization and the role of newcomers and informal leaders in changing group norms (Ashford & Black, 1996; Cooper-Thomas & Anderson, 2006; Cable, Gino, & Staats, 2013).

Taken together, the patterns documented in Appendix A, A1 through A3 indicate that ONA can inform adoption trajectories, attrition and continuity risk (e.g., dependence on a small set of brokers or trusted ties), and coordination bottlenecks, by making knowledge flow, responsiveness, and relational reliability measurable rather than assumed.

Informal networks form an operational layer of knowledge flow that can be measured, interpreted, and influenced. When network maturity, psychological safety, and voluntary collaboration capacity are strong, organizations move faster, sustain learning more consistently, and show measurable performance gains. These findings provide empirical support for treating people-side KM as core infrastructure and set up the argument in Part 5 that enterprise AI requires an integrated socio-technical layer connecting behavioral insights, knowledge quality/provenance, and reinforcement in the flow of work.

Part 5. Implications for this current study

Across Parts 1–3, the literature converges on three core insights that shape enterprise AI adoption.

- First, knowledge quality and provenance influence whether employees view AI-generated recommendations as trustworthy and actionable; technical accuracy is intertwined with consistent curation, metadata standards, and transparent information lineage (Davenport & Prusak, 1998; Nonaka & Takeuchi, 1995).
- Second, technology adoption is fundamentally a social diffusion process shaped by influence networks, trust pathways, and informal collaboration structures, which determine whether new tools spread, stabilize, or stall (Cross & Parker, 2004; Rogers, 2003; Valente, 2010).

- Third, behavioral change occurs through reinforcement in the flow of work; learning and adaptation arise less from episodic training than from repeated cues embedded in daily interactions and collaboration rhythms (Hallsworth & Kirkman, 2020; Eraut, 2004).

Despite broad recognition of these dynamics, enterprise AI systems still struggle because human and social mechanisms remain disconnected from formal knowledge infrastructure. Enterprise platforms such as Microsoft 365, Salesforce, and Workday are optimized for documentation and automation but provide limited visibility into how knowledge actually moves across people, workflows, and influence structures. Feedback tools supply periodic signals, yet these snapshots are weak proxies for the continuous learning loops required for responsible and effective AI integration (Davenport & Prusak, 1998). In the RQ1 (Curation and Governance) arena, this gap affects not only adoption. Governance of knowledge, including behavioral and network knowledge, also improves the quality of what is incorporated into GenAI and agentic AI systems by strengthening provenance, reducing fragmentation, and making retrieval and downstream decisions more reliable.What is missing is an organizational layer that links social dynamics, knowledge quality, and behavioral reinforcement to the information environments where AI learns and operates.

Together, these bodies of work highlight important foundations but leave a clear gap. Existing theories describe the importance of human–system interaction, social structure, and semantic/metadata order (e.g., taxonomies/ontologies and provenance standards), yet they do not offer an integrated architecture that connects semantic governance, behavioral evidence, knowledge provenance, and workflow activity into a continuous learning system. In other words, current KM and ONA scholarship explains *what matters* but offers limited guidance on *how these elements can be operationalized together* in real organizational environments.

Bridging this gap requires a socio-technical framework that links curation and governance with behavioral insight to align human and machine learning loops (Davenport & Prusak, 1998; Nonaka & Takeuchi, 1995). ONA, behavioral analytics, and influence mapping make these dynamics observable and actionable, shifting KM from static repositories toward adaptive organizational intelligence (Cross & Parker, 2004; Valente, 2010). Together, these gaps in governance, behavior, and learning motivate the research questions that guide this study.

Research Questions

Given the conclusions from the literature review, this study addresses a persistent gap in digital transformation and AI adoption: organizations frequently invest in advanced technologies without accounting for the human and social systems that determine whether these technologies take root. Prior work in behavioral science, KM, and network research shows that technology adoption depends not only on tool quality but also on how knowledge flows, how trust and provenance are established, and how autonomic behavioral responses shape employees' willingness to engage with uncertainty (Thaler & Sunstein, 2008; Kahneman, 2011;

Cross & Parker, 2004). When these dynamics are neglected, organizations may experience slower AI uptake, fragmented learning, inconsistent execution, and elevated risk aversion.

Building on these insights, this study examines the effectiveness of three people-side KM interventions for enabling productive and trustworthy AI adoption.

Because this study uses a mixed-evidence design (interviews, systematic desk research, and enterprise behavioral metadata), the "hypotheses" below are framed as testable propositions rather than statistical causal claims. They are assessed through triangulation and pattern matching across evidence streams, with attention to boundary conditions and rival explanations.

Hypotheses 1. Curation and Governance

RQ1. Foundations

How can information architecture, metadata standards, and governance practices — including semantic/metadata normalization where relevant—be designed, socialized, and measured to improve knowledge transparency, provenance, and trustworthiness in AI-supported environments?

This includes examining how baseline "architectural maturity" (coherent integration across core collaboration, workflow, and knowledge systems; stable identities and access controls; and reliable data pipelines) enables a consistent "single source of truth," supported by master data management and curation workflows, enables reliable AI learning, reduces fragmentation, and strengthens organizational confidence in AI-generated outputs. It also includes governing organization-specific context—such as org structure, decision rights, and observed responses to prior interventions or nudges—as a first-class knowledge asset that agentic systems must rely on safely.

Hypotheses 2. Behavioral Insights

In this study, behavioral insights serve two roles: **(1) accelerating and tailor implementation** by revealing where adoption will stall and why, and **(2) informing the design requirements of agentic AI** by specifying decision participants, bottlenecks, and coordination pathways.

To keep these roles conceptually clean,

- **RQ2.1 is scoped to diagnosis and measurement (behavioral evidence)**,
- while **RQ2.2 is scoped to intervention design and testing (nudges).**

RQ2.1 Evidence

How can behavioral and network data reveal the influencers, bottlenecks, relational biases, and trust patterns that shape AI adoption, and how can these signals be governed as organization-specific knowledge that agentic workflows must rely on safely?

This subsection is scoped to observation and inference: identifying behavioral mechanisms and specifying agentic design requirements, not prescribing interventions. In this study, evidence from Organizational Network Analysis (ONA), sentiment assessment, and behavioral analytics serves two functions: (1) implementation diagnosis (where adoption will stall and why) and (2) agentic AI design requirements (who must validate, coordinate, or be included in decisions; where escalation and handoffs occur). These methods are used to detect behavioral drift, influence direction and strength, and psychological safety signals at individual and group levels. In this context, behavioral drift refers to gradual changes in communication, responsiveness, and collaboration patterns relative to an established baseline and is value neutral unless it diverges from intended workflow design assumptions.

This also includes capturing and governing organizational context such as network structure, role-based coordination pathways, observed responses to prior adoption initiatives (including nudges), since these conditions shape both implementation strategy and agentic system behavior.

RQ2.2 Interventions– Micro-Level Behavioral Nudges

How can targeted, evidence-based nudges be designed and tested to strengthen AI literacy, trust, confidence, and responsible use across teams?
This subsection is scoped to action: designing and testing targeted nudges based on the diagnostic signals from RQ2.1. These interventions focus on ***specific, short-horizon*** *behavioral challenges*—for example, when employees show low adoption of a specific tool such as Copilot. In such cases, organizations may deploy localized actions such as reminder cues, role-model demonstrations, bias-reducing framings, or lightweight experiential prompts. The aim is to shift a *specific* behavior in a *specific* context, not to transform the organization's broader learning system.

This level of intervention targets department-specific behavioral barriers by addressing unconscious biases, early misperceptions, and local peer-influence dynamics that shape how employees initially interpret and engage with AI tools. These nudges are designed to shift near-term adoption within a defined team or workflow, but they are not designed to create or sustain organization-wide learning mechanisms.

Hypotheses 3. System-wide Learning Dynamics of Socio-Technical AI Adoption

RQ3. Systems-Level Learning Dynamics

How can people-side KM insights be translated into continuous learning systems that support ongoing organizational readiness for AI?

Whereas RQ2.2 focuses on short-term, department-level behavioral shifts, RQ3 examines **how organizations develop nonstop, ongoing, system-wide learning capabilities**.This line of inquiry examines how organizations build *persistent socio-technical learning capabilities*—mechanisms that help them detect, interpret, and respond to barriers to AI adoption across roles, teams, and workflows over time. The focus is not on encouraging use of a single tool but on creating adaptive systems that monitor collective behavioral evidence, strengthen knowledge flows, and integrate feedback loops into operational and decision making practices, informed by business intelligence. These systems allow organizations to continuously recalibrate norms, processes, and learning pathways even as specific AI tools change, ensuring sustained readiness and accountable use of emerging technologies.

The following section outlines the methodological approach used to investigate these questions.

Methodology

The approach uses mixed methods that integrate practitioner interviews, public executive commentary, and large-scale behavioral data from enterprise knowledge-sharing environments. The study's logic model is below. (The "Organizational Intelligence Loop" [OIL], a response to our findings, will be introduced in the Discussion below.)

Category	Theme	Practices / Evidence in This Study	Design Implications (Framework Output)
AI Adoption Challenges	Fragmented knowledge; uneven adoption; lack of trust	Barriers identified in practitioner interviews; public commentary on trust and governance; patterns from LEAD enterprise data	Establishes the motivation for a socio-technical adoption model
Knowledge & Behavioral Foundations	People-side KM, tacit knowledge, network influence, psychological safety	Literature review synthesis; mapping constructs (trust, influence, metadata, reciprocity)	Provides the theoretical foundation for OIL's People, Information, and Process layers
Research & Evidence Generation	Mixed-method insights across enterprise contexts	Interviews, commentary, organizational behavior data; development of hypotheses H1–H3 (RQ1- RQ3)	Produces analytical structure to interpret adoption dynamics
Socio-Technical Findings	Behavior, Trust, and Learning Dynamics	Identification of trust pathways, informal influencers, metadata gaps, behavioral nudges, learning conditions	Specifies the mechanisms that the OIL framework must operationalize
OIL Framework Architecture	Integration of human, information, and workflow signals	Development of four OIL layers: People, Information, Process, Agentic AI	Provides an adaptive learning architecture for enterprise AI adoption
Strategic & Future Implications	Organizational design; multi-agent governance; accountability	Considerations for scaling agentic AI, trust calibration, and continuous system alignment	Situates OIL within broader questions of enterprise AI governance, system alignment, and future socio-technical design

1. Research Design

This study employs a multi-method qualitative design integrating four complementary approaches:

- Literature review,
- Semi-structured practitioner interviews,
- Systematic desk research of public commentary by executives, transformation leaders, thought leaders, and researchers.

This design aims to connect practitioner experience and industry discourse to understand Agentic AI adoption and inform emerging frameworks for people-centered AI adoption.

2. Part 1 Interviews

Practitioner Perspectives

Semi-structured interviews were conducted with seven senior practitioners, with experience in customer success, marketing, program management, software development, product

management, higher education and knowledge management. They are representing distinct AI adoption contexts:

- **Christina Libs — Program Manager & Knowledge-Management Entrepreneur, [Company Redacted]**
  A senior leader in knowledge management and customer experience, specializing in self-service systems, digital support, and AI-driven service operations.
- **Dr. Edward J. Hoffman — Former Chief Knowledge Officer, NASA; Founder & CEO, Knowledge Strategies LLC; Senior Lecturer, Columbia University**
  Founding Director of NASA's Academy of Program/Project and Engineering Leadership (APPEL) and author of *The Smart Mission: NASA's Lessons for Managing Knowledge, People, and Projects*. He teaches knowledge management and organizational learning in Columbia University's Information & Knowledge Strategy (IKNS) program.
- **Eve Porter-Zuckerman — Independent Consultant; Former Chief Knowledge Officer, Isaacson, Miller**
  Former CKO at a premier executive recruiting firm and consultant specializing in organizational change and professional-services knowledge management. Eve teaches AI and semantic reasoning at Columbia IKNS.
- **Mihai Pohontu — Founder & CEO, Amber Studios**
  Leads a 900-person creative studio building GenAI tools for internal workflows and top-tier "AAA" gaming clients—games produced with blockbuster-level budgets, teams, and production quality.
- **Nabeel Ahmad — Co-Founder & CSO, ChangeForce; Adjunct Professor, Columbia University**
  Adjunct Professor of Human Capital Management at Columbia University and founder of *Change.AI*; former Head of Market Development & Industry Learning at KPMG.
- **[Name Redacted] — Tech Lead, Twilio**
  Leads internal generative-AI enablement for Customer Success.
- **[Name Redacted]— Y Combinator (YC)-Affiliated Founder and Y Combinator Partner, [Company Name Redacted]**
  A senior technology founder leading a late-stage, venture-backed digital platform company with over 500 employees.

Interviews lasted 30-45 minutes and followed an open, semi-structured guide organized around four core themes:

- **Barriers and enablers of AI adoption** — *What slows or accelerates the use of new AI tools, such as internal knowledge assistants for engineering, customer support, or sales teams?*

- **Influence networks** — *Who actually shapes employees' decisions to adopt or reject AI tools, beyond formal trainers or leaders?*

- **Behavioral nudges and cultural signals** — *Do you use incentives, default values (choice architectures), social norms, or hands-on seminars to encourage experimentation and adoption?*

- **AI readiness and literacy** — *How do you help employees gain enough understanding and confidence to use AI responsibly, and how do you build AI literacy beyond basic*

training? Before asking this question, the researcher clarified that AI literacy in this study extends beyond tool use to include how AI fits into existing workflows, how people-side dynamics influence adoption, and how AI connects to broader business and cultural practices. This clarification was intended to keep the discussion on workplace use conditions rather than unit pricing, technology maturity, or other commercial and market considerations that fall outside the study's focus.

Ethical Considerations:

The full question set is provided in the appendix. Interviews were designed to allow participants to elaborate on unexpected factors and provide concrete organizational examples. All sessions were recorded or documented with consent. Participants could choose anonymity or attribution. Quotes were anonymized and lightly edited for clarity. To protect confidentiality, interview recordings and transcripts were not shared outside of the research process itself, and will be destroyed at the end of the research. Participants provided informed consent and could withdraw at any time.

3. Part 2 External Public Commentary

To complement practitioner data and situate findings within the broader industry discourse, the study systematically collected and coded public commentary (2023–2025) from executives, transformation leaders, thought leaders and researchers. Public sources included LinkedIn, X (Twitter), professional blogs, and industry interviews, with sources limited to verified accounts and recognized experts or publications.

- **Search scope and process:** Searches combined key terms such as *"AI adoption," "knowledge management," "organizational network analysis," "behavioral nudges," "trust," "AI governance."* Queries did not require sources to contain all terms; results were included when they addressed one or more of these themes in the context of enterprise AI, with the search tool also permitting closely related synonyms. Although behavioral concepts are relevant, the search did not rely heavily on formal organizational-psychology terminology. Instead, it focused on how business leaders actually talk about these issues in practice, as they build or adopt agentic AI. We were particularly interested in their use of people-side KM despite not having a background in organizational behavior or knowledge management. Narrower terms such as "social proof" and "defaults" were included to capture related synonyms, were included to capture adjacent patterns without constraining retrieval to a fixed keyword list.
- **Search process**: The search process used multiple tools, including Google Search, Grok, Perplexity, and ChatGPT, to cross-reference publicly available materials. For each search, named executives and thought leaders were provided as anchors, using their verified LinkedIn or X (Twitter) profiles to locate related publications, interviews, posts, or talks that referenced the above keywords or addressed challenges in enterprise AI and GenAI adoption. Only content that could be independently verified or attributed to credible professional sources was retained for analysis.

- **Inclusion criteria:** Posts had to express an opinion on whether *people-side knowledge management* (as defined in the research questions) helps or hinders AI adoption. Both supportive and skeptical perspectives were retained.
- **Sampling approach:** The researcher set a minimum target of at least 30 distinct public commentators and prioritized sources published within the most recent five years when available. The final sample included 32 sources.

Findings

In this section we review the evidence from each of the research sources.

1. Findings from Semi-Structured Practitioner Interviews and Systematic Desk Research of Public Executives

1.1 Insights from Interviewees

Table 1 consolidates the practitioner interviews by showing (1) which research questions each interviewee most directly addressed (ranked by their dominant focus) and (2) the main themes they emphasized when describing barriers and enablers of enterprise AI adoption. Rather than treating interview input as one undifferentiated set of anecdotes, the table makes visible how different practitioner roles naturally gravitate toward different explanatory layers: governance and traceability (RQ1), behavioral and network evidence (RQ2.1), micro-interventions and nudges (RQ2.2), or system-wide learning dynamics and workflow integration (RQ3).

Table 1 Interviewee Focus Areas

Interviewee Focus Areas: Corresponding Research Questions & Main Themes		
Interviewee	**Research Questions Addressed (Ranked by Their Dominant Focus Area)**	**Main Themes They Emphasized (Ranked by Their Dominant Focus Area)**
Christina Libs	RQ2.1 (behavioral evidence: influence, trust patterns) RQ2.2 (behavioral interventions: peer demos, social proof) RQ3 (learning through experimentation) Partial RQ1 (governance gaps)	• People & Behavior as Core to AI Adoption • Informal Networks & Trust • Behavioral Micro-Barriers (fear, norms) • Workflow & Process Integration • Learning & Reskilling
Dr. Edward J. Hoffman	RQ1 (governance, problem definition, traceability) RQ2.1 (influence mapping, behavioral bottlenecks) RQ2.2 (reflection as nudge) RQ3 (continuous learning systems, cross-functional loops)	• Curation & Governance • Knowledge Architecture / Context Design • Informal Networks & Influence Bottlenecks • Behavioral Micro-Barriers (politics, misalignment) • Learning & Reskilling (reflection loops) • Organizational Change Velocity
Eve Porter-Zuckerman	RQ1 (safety, governance clarity) RQ2.1 (influence, fear, culture) RQ2.2 (peer champions, safe-to-fail practice) RQ3 (iterative learning and workflow integration)	• Ethical & Responsible Analytics • Curation & Governance (safety, guardrails) • Behavioral Micro-Barriers (fear, norms) • Informal Networks & Trust (peer champions) • Learning & Reskilling (micro-practice, workflow fit)
Mihai Pohontu	RQ2.1 (informal networks, influence spread) RQ2.2 (nudges through social proof, persistent exposure) RQ3 (learning via real projects, iterative workflows) Partial RQ1 (governance concerns for IP/copyright)	• Informal Networks & Trust (grassroots adoption) • Organizational Change Velocity (artists' resistance) • Workflow & Process Integration (top-down fails) • Behavioral Micro-Barriers (fear, identity threat) • Learning & Reskilling (internal projects)
Nabeel Ahmad	RQ2.1 (behavioral evidence, resistance patterns) RQ2.2 (workflow-based interventions, experiential learning) RQ3 (readiness, literacy in context) Partial RQ1 (governance maturity gaps)	• People & Behavior as Core to AI Adoption • Behavioral Micro-Barriers ("AI shame," habits) • Organizational Change Velocity (risk aversion) • Learning & Reskilling (sandboxing, workflow fit)
Twilio Tech Lead	RQ2.1 (informal influence, peer modeling) RQ2.2 (subtle nudges, visibility) RQ3 (learning-by-doing) Partial RQ1 (constraints from compliance/governance)	• Informal Networks & Trust (IC-driven adoption) • Workflow & Process Integration • Behavioral Micro-Barriers (uneven agency) • Learning & Reskilling (sandbox practice)
YC-Affiliated Founder	RQ2.1 (influence from motivated ICs) RQ3 (peer-driven learning, emergent literacy) Minimal RQ1 or RQ2.2 but still behavior-aware	• Informal Networks & Trust (ground-up adoption) • Knowledge Architecture / Context (IC-created guidance) • Workflow & Process Integration (embedding AI in team processes)

(Source: "OIL Interview Candidate List – Name List CSV," coded by author. Name & brief intros are listed in Appendix C.)

1.2 Findings from Systematic Desk Research of Public Commentators (Executives, Transformation Leaders, Thought leaders, Researchers)

Table 2 synthesizes the desk-research sample by mapping each research question to (a) recurring themes in the public and institutional commentary reviewed, (b) representative stakeholder roles and source types that most often emphasized those themes, and (c) the stated rationale for why each issue matters for enterprise AI adoption. This table is a structured synthesis of perspectives and does not represent a frequency count or statistical weighting.

Table 2. Summary of Research Questions, Themes, Roles, and Core Rationale

Summary of Research Questions, Themes, Stakeholder Roles, and Core Rationale			
Research Question	**(a) Themes**	**(b) Who Mentions It**	**(c) Why These Roles Emphasize It**
RQ1. Curation & Governance — How information architecture and governance improve transparency, provenance, trust.	• Knowledge Architecture & Context Design • Workflow & Process Integration	AI CEOs; Network scientists; People-analytics leaders; MIT/ServiceNow; BIT; NHS	These roles focus on system-level reliability, metadata quality, and workflow coherence—foundational conditions for trustworthy AI outputs and a single source of truth.
RQ2.1 Behavioral Insights — Evidence How behavioral and network data reveal influence, bottlenecks, trust patterns.	• Informal Networks & Hidden Influence • People & Behavior as Core to Adoption • Organizational Change Dynamics	People-analytics leaders; ONA researchers; Organizational psychologists; Future-of-work strategists; Transformation experts	Their work centers on uncovering tacit influence, trust dynamics, and behavioral patterns that explain uneven adoption across teams.
RQ2.2 Behavioral Insights — Interventions How targeted nudges strengthen literacy, trust, responsible use.	• Learning & Reskilling • Behavioral Micro-Barriers	Educators (MIT/Wharton); DeepMind; Workforce-intelligence leaders; BIT/behavioral scientists	These roles design localized interventions and literacy pathways to shift specific behaviors, reduce friction, and build confidence in AI tools.
RQ3. Learning Dynamics of Socio-Technical AI Adoption — How to build continuous organizational readiness systems.	• Adaptive Knowledge Architecture • Iterative Workflow Integration • Change Readiness Monitoring • Evolving Influence Networks	AI CEOs; Applied AI leaders; Transformation experts; ONA researchers; Network scientists	These roles view AI adoption as an ongoing socio-technical learning process requiring dynamic sensing, recalibration, evolving workflows, and monitoring of shifting collaboration patterns.

(Source: “OIL Interview Candidate List – Web Resource CSV,” coded by author. Original links and summaries are provided in Appendix C.)

Table 3. Representative Public Commentators

Table 3 highlights well-recognized public voices in enterprise AI and workplace transformation, included because their commentary is widely circulated and influential in shaping how organizations interpret adoption, governance, and readiness. The table shows how these high-signal, non-random sources map to the study’s research questions and themes.

Person	Affiliation / Role	Key Insights	Primary Research Alignment (Ranked by Their Dominant Focus Area)
Aaron Levie	CEO, Box	AI performance is constrained by missing organizational context. Tacit knowledge must be made explicit through metadata, documentation, and structured workflows. Teams must develop new habits of disciplined knowledge hygiene.	RQ1: Metadata, provenance, workflow clarity. RQ3 (secondary): Reskilling toward structured knowledge behaviors.
Adam Grant	Professor, Wharton School	Effective AI adoption depends on human adaptability, unlearning old norms, and strengthening collaboration and judgment. Culture change and mindset shifts are essential.	RQ2.1: Social norms, role modeling, behavioral patterns. RQ2.2: Reframing, unlearning as behavioral interventions. RQ3: Continuous learning agility.
Aditya Challapally	Researcher, MIT / NANDA Lab	Successful AI pilots rely on behavioral design—collaboration contracts, trust verification, escalation paths—not superior algorithms. Failures stem from unclear workflows and weak human-centered processes.	RQ2.1: Trust dynamics, collaboration norms. RQ2.2: Behavioral design as structured interventions. RQ1 (secondary): Need for workflow clarity and explicit context.
Al Adamsen	Founder & Director, Future of Work Project	AI adoption requires human systems design, cross-functional governance, and alignment of psychological needs. Leaders must create adaptive, confidence-building environments.	RQ1: Governance alignment. RQ2.1: Psychological safety and influence pathways. RQ3: Adaptive learning and sensemaking.
Alexis Fink	VP, People Analytics, Meta	AI success comes from understanding work at the task level, redesigning roles, and aligning analytics to business value rather than dashboards. Blending human-centered design with structured work analysis improves readiness.	RQ1: Task/workflow definition and governance clarity. RQ3: Cultural learning and organizational sensemaking. RQ2.1 (secondary): Cultural signaling from leadership.
Andrew Ng	Co-founder, Google Brain & Coursera	AI fluency emerges through experiential learning—"learning to drive by driving." Hands-on experimentation enables skill development faster than structured instruction alone.	RQ2.1: Priming and defaults through real-task interaction. RQ3: Continuous reskilling and iterative learning loops.
Cole Napper	Head of Workforce Analytics, Lightcast	People analytics and GenAI reshape workforce decisions only when paired with clean data, clear workflows, and governance. Storytelling and embedded AI usage shape norms and accelerate adoption.	RQ1: Data maturity and structured workflows. RQ2.2: Storytelling and social proof & Embedded AI as default choice architecture.
Dawn Klinghoffer	VP, HR Business Insights, Microsoft	infrastructure, embedding insights into decisions while maintaining transparency and ethical data use.	RQ1: Responsible analytics governance. RQ3: Enterprise learning and insight loops.
Helen Bevan	Chief Transformation Officer, NHS Horizons	Tacit knowledge and informal networks—not explicit knowledge—drive innovation. Psychological safety and relational trust determine whether expertise flows across boundaries.	RQ2.1: Informal influence, trust networks, tacit knowledge. RQ3: Decentralized learning and reflective practice.
Jensen Huang	CEO, NVIDIA	AI-human collaboration requires contextual understanding, ethical use, and continuous mutual learning. Human judgment shapes how systems evolve.	RQ1: Need for coherent context architecture. RQ2.1: Human judgment signals as behavioral guidance. RQ3: Continuous human-AI co-learning.

(Source: "OIL Interview Candidate List – Web Resource CSV," coded by author. Original links and summaries are provided in Appendix C.)

1.3 Role-Based Overview of Themes Across Interviews and Public Commentary

Across the 39 sources examined in this study, including seven practitioner interviews and 32 public commentaries, a consistent insight emerges: the success of AI adoption depends not on technology alone but on the interaction between architectural maturity and human capability. "Architectural maturity," as used here, refers to the basic reliability and integration of the systems that underpin organizational work, including communication platforms, workflow tools,

information repositories, and the data pipelines that feed AI models with coherent, traceable inputs. Although not stated explicitly in the research questions, architectural maturity forms the substrate that makes governance, behavioral insights, interventions, and continuous learning possible. Interviewees and public commentators claimed that without stable systems, organizations lack the visibility and data required to address the very challenges posed in RQ1 through RQ3. Interviewees reinforced this implicit dependency on architectural maturity, describing how fragmentation, shadow AI, unclear guardrails, and inconsistent information flows hinder even highly capable employees from adopting AI effectively.

Architectural maturity is necessary, but not sufficient: Interviewees and public commentators asserted that psychological safety, peer influence, and literacy depend on additional systems that make behaviors observable and knowledge flows transparent.Specifically, both interviews and public commentaries indicate this dual need for AI investments translate into measurable value:

1. Trust, influence networks, and ongoing learning
2. Reliable outputs enabled by robust metadata and a well-managed underlying knowledge base.

Technically mature environments cannot succeed without shared understanding, trusted relationships, and social norms that make experimentation safe. Across sources, AI adoption emerges as a mutually reinforcing socio-technical process in which human and architectural conditions shape each other.

Within the 32 public commentaries, distinct role-based patterns appear in how experts describe the challenges of AI adoption:

- People analytics leaders, network scientists, and transformation practitioners in complex systems such as health care emphasize informal networks, tacit knowledge, collaboration bottlenecks, and hidden influence. Their vantage point leads them to see social structures as foundational to AI readiness.
- System-level leaders, including AI founders and large-company executives, foreground governance, transparency, and responsible analytics, reflecting accountability concerns.
- Applied AI practitioners at MIT, ServiceNow, and Insight222 describe workflow and process integration as central, observing that AI succeeds only when embedded into real work.
- Technical CEOs, DeepMind researchers, ONA scholars, and behavioral insights experts highlight metadata quality, knowledge architecture, and contextual models as the primary constraints on AI performance.

Taken together, these perspectives reveal that most contributors speak deeply to only one dimension of the socio-technical system, shaped by the scope of their professional responsibilities. Only a small subset of commentators, largely people analytics leaders and future-of-work strategists, consistently integrate behavior, governance, workflow, and architecture.

Technical founders commenting publicly focus on architecture because that is where their responsibilities lie, while practitioners closer to implementation discuss behavioral and relational dynamics in more detail because they encounter them directly. Behavioral scientists focus on

psychological mechanisms, while knowledge managers attend to trust and literacy; each group contributes a partial view of a larger system.

The practitioner interviews offer additional insights that rarely surface in public commentary. Interviewees described emotional resistance, uncertainty, and role-specific anxieties, phenomena not typically addressed in high-level narratives.

They also discussed micro-level workflow friction, local politics, influence asymmetries, and gaps between conceptual literacy and applied competence. They also highlighted fragmented ownership, unclear safety boundaries, and governance misalignment as practical barriers to adopting AI in day-to-day work. These differences do not contradict public commentary but reveal how role and context shape what individuals are positioned to observe.

Across all 39 sources, the collective findings demonstrate that AI adoption is a multi-layered socio-technical learning process shaped by governance clarity, behavioral mechanisms, informal influence networks, workflow context, and iterative feedback loops. Architectural maturity alone cannot ensure adoption, nor can behavioral interventions succeed without supportive structures. Understanding AI adoption therefore requires examining how these elements interact across roles, systems, and moments of practice. The following analysis turns to the study's four research questions to explore these interactions in greater depth and identify the challenges and opportunities that emerge from this combined dataset.

2. Findings Organized by Research Hypotheses

2.1 Curation and Governance (RQ1): How Information Architecture and Metadata Transparency Build Trust in AI

Across interviews, participants noted that trust in AI depends less on technical sophistication and more on clarity about where information comes from, how it is maintained, and how it flows across the organization. Consistent information architecture, supported by shared definitions, metadata standards, and governance processes, was repeatedly identified as a prerequisite for reliable AI outputs. When employees cannot see or understand the knowledge base that AI relies on, they question not only individual answers but the system as a whole.

- **Christina Libs** observed that employees often do not know how AI works or where its information comes from. She emphasized that reliability depends less on algorithms and more on the system's "diet," referring to the quality, structure, and curation of the underlying knowledge base. When documentation is scattered across chat logs, tickets, and private folders, employees encounter outputs that feel uneven or arbitrary. As she explained, "When you don't trust the inputs, you can't trust the answers." Her experience illustrates how fragmentation in knowledge sources and metadata produces fragmentation in user confidence.
- The **Twilio Tech Lead** highlighted a related but distinct issue: governance uncertainty. Privacy and compliance policies had not kept pace with emerging AI workflows, leading

employees to default to restrictive or covert experimentation. Instead of clarifying what data could be used and how, organizations relied on caution, which created uneven adoption and uncertainty about acceptable boundaries. His account shows that trust requires visibility not only into data sources but also into the rules and constraints that govern their use.

- **Nabeel Ahmad** observed that when organizational systems feel outdated or do not offer sanctioned tools that meet employee needs, people often "go rogue." Employees turn to unsanctioned AI tools out of necessity rather than defiance, reflecting low trust in existing systems and uncertainty about ownership or reliability. He noted that many leaders experience "AI shame," avoiding use or discussion of AI for fear of misunderstanding it. Nabeel argued this created an authority vacuum. In his view, "rogue" behaviors emerge when employees lack confidence in what is approved and authoritative, including which sources are trusted, what rules apply, and what boundaries govern acceptable use, signaling gaps in clarity, communication, and governance.
- In creative industries, **Mihai Pohontu** concurred with Ahmad. Mihal described how uncertainty about data lineage and content ownership create hesitation and slow experimentation. As he explained, partners are not afraid of creativity but of "not knowing who owns what." Without transparent metadata tracing and clear governance, even internal AI-generated assets became sources of risk. Pohontu emphasized that innovation stalls when provenance and ownership are ambiguous, as employees do not want to be the first to make a mistake.
- Even startups face this pattern. A **YCombinator partner** explains that teams adopt tools like Claude and Copilot organically, but leadership still has to ensure that sensitive data does not leak. Technical expertise does not remove the risk; it heightens the anxiety around it. "Everyone's excited," he says, "but no one wants to be the one who crosses the line."
- **Dr. Edward J. Hoffman** emphasized the systemic nature of these trust challenges. He noted that departments often operate with different definitions, processes, and governance assumptions, which creates inconsistencies in how information is interpreted and used. When collaboration is fragmented and information flows are disconnected, decisions lose traceability, and AI systems built on top of these foundations inherit the same inconsistencies. Hoffman explained that trustworthy AI depends on an integrated knowledge architecture in which leadership intent, curated data sources, and workflow context are aligned. Such coherence enables employees to understand how inputs connect to outcomes and to trace how information supports the decisions they see.
- **Public commentary** reinforced these insights. **Aaron Levie** stressed that "fragmented data systems create fragmented organizations," while **Alexis Fink** described metadata transparency as the "invisible foundation of trust." Both perspectives align with findings from the Behavioural Insights Team that transparency reduces uncertainty, increasing adoption when employees understand where information comes from and how it is updated.

Taken together, these findings support RQ1 and extend the cross-source synthesis in Finding 1.3: AI adoption depends on "architectural maturity," or well-defined information architecture and governance processes that make metadata, provenance, and knowledge flows visible to employees. A consistent "single source of truth," supported by master data management, curation workflows, and clear usage boundaries, enables traceability, reduces fragmentation, and strengthens organizational confidence in AI-generated outputs.

2.2 Behavioral Evidence (RQ2.1): Influence Networks and Trust Patterns Shape Adoption Trajectories

Across interviews, participants emphasized that AI adoption is shaped primarily by individual's biases and heuristics and groups' informal social dynamics rather than only formal training or formal authority. Although interviewees did not use terms such as Organizational Network Analysis (ONA) or behavioral analytics, the influence patterns they described map directly onto the kinds of relational signals these methods can capture. Their accounts highlight mechanisms such as messenger effects, perceived norms, identity, and psychological safety that determine how employees interpret early experimentation.

- At Amber Studios, **Mihai Pohontu** observed that AI experimentation accelerated once early adopters shared their results across teams. Their visible examples reduced uncertainty and signaled that experimentation was acceptable and valuable.
- The **Twilio Tech Lead** described a similar dynamic in which small peer groups generated momentum that formal programs could not replicate. The Twilio Tech Lead posited that these behaviors illustrate how informal influence clusters and local norms shape adoption trajectories, and how such clusters can be identified through network-based indicators in practice.
- In more traditional environments, influence pathways often transmitted caution rather than curiosity. **Christina** Libs explained that employees hesitated to engage with AI when they did not understand how tools worked or what information they relied on. She observed that employees looked to trusted peers, not managers, to determine whether AI use felt safe or relevant. Libs posited that these patterns reflect the underlying trust networks, psychological safety signals, and informal opinion leaders that behavioral and relational data can reveal.
- **Dr. Edward J. Hoffman** noted that leaders often lack visibility into who employees rely on for interpretation and reassurance. He described how executives, IT teams, and external experts frequently operate in silos, creating bottlenecks that slow experimentation and misalign decision-making. Hoffman claimed that these bottlenecks and influence asymmetries are examples of structural patterns that can be surfaced through network and behavioral analysis.
- **Public commentary** reinforces these themes. **Helen Bevan** emphasized the importance of tacit knowledge and relational trust in complex systems, while **Adam Grant** highlighted mindset, identity, and social contagion as core mechanisms of organizational learning.

Collectively, these insights support RQ2.1: AI adoption depends on mapping and understanding informal influence networks and trust patterns. As we saw in the literature review, behavioral and network-based data provide the empirical means to detect influencers, bottlenecks, relational biases, and emerging psychological safety trends that shape whether early engagement accelerates or stalls.

2.3 Behavioral Interventions (RQ2.2): Embedding Behavioral Nudges in Daily Workflows to Strengthen AI Readiness

Interviewees described how targeted, evidence-based nudges can strengthen AI literacy, trust, and responsible use by addressing discrete behavioral barriers that emerge during early experimentation. Rather than attempting to reshape the entire learning system, these interventions focus on specific moments when employees hesitate to adopt or correctly apply an AI tool. They operate through well-established behavioral mechanisms such as choice architectures (e.g., default values), salience, ease, affect, and timely prompts, which help individuals reinterpret AI use as safe, useful, and aligned with their workflow.

- These interventions occurred within the flow of work, not as detached training modules. At **Twilio**, employees began adopting AI tools once they were embedded directly into communication and workflow platforms. While the interview did not specify the structure of formal training, the **Tech Lead** emphasized that adoption accelerated when employees could experiment with AI in the context of their actual tasks. He suggested that embedding AI into existing routines reduced friction, made the tools more immediately useful, and allowed teams to learn through direct interaction rather than relying solely on standalone instructional sessions.

- Practitioners also emphasized lightweight, context-specific nudges that surface relevance and build confidence. At **Amber Studios**, **Mihai Pohontu** described how short weekly stories of creative breakthroughs increased positive salience and helped employees reinterpret AI as a supportive tool rather than a risk. **Eve Porter-Zuckerman** similarly highlighted the value of micro-experiments, small and low-stakes actions such as refining a message, which allow employees to experience early wins. She posited that these interventions rely on affect and simplicity, making success emotionally resonant and encouraging continued engagement.

- **Nabeel Ahmad** observed that employees adopt AI more readily when its use is tied directly to daily tasks instead of general training programs. Ahmad claimed that adoption improves when tools address clear pain points and produce tangible time savings, supported by safe environments in which employees can experiment within their real work context. These workflow-integrated cues shift behavior not through persuasion but by making AI easier to apply to familiar tasks.

- **Public commentary** echoed these themes. **Andrew Ng** emphasized learning by doing as a form of behavioral priming that builds confidence through repeated interaction. **Cole Napper** noted that embedding AI features directly into workforce systems establishes natural defaults that shape habits over time.

Together, these insights demonstrate how targeted, evidence-based nudges such as embedded prompts, role-model demonstrations, emotionally resonant examples, and micro-experiments can be designed and tested to address specific adoption barriers. Consistent with Hallsworth & Kirkman (2020), these interventions improve literacy, trust, and confidence by aligning behavioral cues with actual work practices and by reducing early misperceptions, uncertainty, and other common adoption frictions. In this way, localized nudges help employees reinterpret

AI tools as safe, relevant, and useful, enabling responsible use long before broader cultural or systemic changes take hold.

2.4 Learning Dynamics (RQ3): Building Continuous Socio-Technical Systems for Sustained AI Readiness

Across interviews and public commentary, a consistent insight emerged: AI readiness is not a static skill but a continuously reinforced organizational capacity. Participants described how employees learn most effectively when they can observe real outcomes, adjust their approaches, and receive ongoing guidance. This reflects RQ3's proposition that readiness develops through socio-technical systems that support interpretation, reflection, and adaptation over time.

Interviewees in fast-moving environments described learning as an iterative, evidence-driven process. **The YC-affiliated founder** noted that teams refined their prompting strategies by experimenting, sharing observations, and rapidly integrating improvements into workflows. He contended that such cycles illustrate how continuous feedback loops—rather than one-time instruction—enable teams to calibrate their practices in response to evolving needs.

In more traditional settings, interviewees highlighted barriers that prevented these adaptive loops from forming: Lack of goal post, experimentation-permission, measurement, isolation.

- **Christina Libs** observed that employees hesitated to experiment when they lacked understanding of how AI systems worked or how they related to daily tasks, reinforcing the need for structures that clarify purpose, expectations, and boundaries.
- **Nabeel Ahmad** added that fear, uncertainty, and low confidence deterred employees from engaging deeply with AI, suggesting that readiness depends not only on capability but on environments that enable safe, contextualized experimentation.
- **Dr. Edward J. Hoffman** emphasized that many organizations lack mechanisms to evaluate whether AI tools truly address defined problems, creating a "literacy–application gap." He argued that organizations require shared frameworks that connect intent, governance, and workflow context so employees can interpret outputs and improve system use over time.
- **Mihai Pohontu**'s account showed that isolation is also a risk. He showed how communities of practice and shared artifacts make learning visible and reinforce collective evolution of skills.

Public commentary made the implications of these gaps in systemic learning explicit.

- **Andrew Ng'**s teaching model assumes that fluency in AI develops through repeated, hands-on interaction with tools, often coupled with guided reflection on what works and why.
- Recent work from **the Behavioural Insights Team** treats AI adoption as a behavioural challenge and recommends iterating interventions based on diagnosed behavioural barriers and evidence from practice.

- **Adam Grant** addressed that sustained learning depends on identity, meaning-making, and social reinforcement.

These perspectives collectively extend the interview data by articulating how organizations can operationalize continuous learning at scale: through feedback loops, readiness monitoring, evolving influence maps, and mechanisms that surface behavioral evidence across teams.

Together, these findings support RQ3: people-side KM enables the design of continuous socio-technical learning systems by revealing where employees struggle, how knowledge flows, and which relational dynamics support or hinder adaptation. These learning systems draw on behavioral evidence and iterative sensemaking, and they are strengthened by the governance clarity and information discipline examined in RQ1. When these elements operate together in organizational routines, leaders can recalibrate norms, processes, and learning pathways as demands shift. In doing so, they ensure that AI "architectural maturity" (see Finding 1.3) evolves alongside organizational context, creating sustained readiness rather than momentary proficiency.

The Findings highlight consistent socio-technical forces that determine whether AI adoption accelerates or stalls. They establish the foundation for the discussion that follows, which interprets these Findings in relation to the study's research questions and broader implications for organizational design and AI integration.

Discussion & Implications

In our discussion section we respond to the research questions (Section 1) and explore implications as a novel model, the Organizational Intelligence Loop (Section 2).

1.1 Integrating Findings and Research Questions

Drawing from a mixed-method analysis of 39 sources, including practitioner interviews, public commentaries, and relevant literature—including prior ONA research and patterns documented across LEAD's multi-organization dataset—the study finds that productive AI adoption depends not on technical sophistication alone but on the alignment of three elements: information governance (RQ1), social and behavioral dynamics (RQ2), and ongoing learning processes (RQ3).

The literature review shows that organizations require a multidimensional set of elements – consistent knowledge architecture, shared language, and psychological safety – to support complex socio-technical change (Cross and Parker, 2004; Davenport and Prusak, 1998; Borgatti et al., 2018). These play out in the Findings:

- Reliable enterprise systems provide an important foundation for this by making coordination and workflow continuity possible. Yet, knowledge architecture and reliability are necessary but not sufficient.
- Social interaction patterns and everyday behavioral habits determine whether this foundation is used effectively.
- Relationships, communication norms, and learning practices shape how people interpret AI tools and how they incorporate them into their work, turning structural potential into practical outcomes.

Across the sources analyzed, several consistent patterns emerged, each corresponding to one of the study's hypotheses. The findings below are organized through RQ1, RQ2, and RQ3.

1.1.1 Curation and Governance — Transparency Builds Trust

Findings from interviews and enterprise data show that employees trust AI when they can trace how it learns, and verify that information is accurate, current, and correctable.

1. **Data Credibility**: When data sources are fragmented or ownership is unclear, confidence erodes, When data and metadata are consistent and accountable, trust grows. Interviewees also noted that both discipline and transparency in data handling shaped their confidence in AI-supported workflows.
2. **Visible Governance**: Interviewees claimed that governance rather than technology emerged as the main perceived differentiator. Many interviewees and public commentary claims that organizations lacked a single source of truth: a unified and visible foundation connecting data provenance, ownership, and updates across systems. In many cases, trust broke down upstream, before implementation, because leaders could not clearly define what problems AI was intended to solve or how information, intent, and ownership connected across departments.
3. **Explicit Alignment**: Building trust, therefore, requires purposeful governance that links data flow to organizational goals and accountability. When aligned, information architecture shifts from a static compliance record to an active system of transparency and shared knowledge, enabling people to trace how decisions and information connect.

These findings address the first research question by showing that the all-important trust component of AI adoption requires data credibility (and purposeful information design), visible governance, and explicit alignment are essential for creating sustainable trust in AI use. They also suggest that, in practice, trust depends not only on metadata management but on the ongoing harmonization of data and metadata across systems, which interviewees described as necessary for maintaining clarity and reliability in a dynamic organizational environment.

1.1.2 Behavioral Insights — Policy, Peers, and the Spread of Confidence

Behavioral and network data show that organizational structure and social dynamics jointly determine how quickly AI is adopted. Interviewees provided numerous behavioral insights that illustrated this relationship. Several noted that while policy establishes boundaries and accountability, most AI policies in this study were either missing or overly restrictive, leaving employees uncertain about how to experiment safely. In these situations, informal influence

networks filled the gap. Interviewees described the need for behavioral insights, including social structures:

1. **Leveraging Norms**: trusted peers modeled responsible use, shared early examples, and created conditions that made engagement feel acceptable.
2. **Showcasing Messengers and Using Salience**: Similar behavioral cues also appeared in the public commentary, where executives emphasized the importance of early role models and visible success stories.
3. **Making Connections visible**: ONA made visible patterns that interviewees had previously observed informally. Interviewees claimed that adoption spread more readily when credible connectors were recognized and supported through pilots or cross-functional work, and that adoption slowed when these individuals were isolated.
4. **Making Nodes Nodes Visible**: Organizations with visibility into internal expertise—referred to in the literature as transactive knowledge and described by interviewees as “knowing who knows what”—were better able to align resources and talent around shared goals instead of running disconnected experiments. ONA does not replace governance; rather, leaders who interpret ONA insights can use them to strengthen coordination and reinforce supportive influence pathways.
5. **Using Influence Architectures Strategically and Ethically:** Interviewees also highlighted specific behavioral nudges that lowered hesitation during early adoption. These included embedding AI capabilities directly into everyday workflow tools to reduce friction, offering short and concrete examples to make benefits immediately understandable, and providing opportunities for low-stakes experimentation. Public commentary echoed these themes, noting that early demonstrations, visible defaults, and well-timed prompts can shift employee perceptions, but require transparency and guardrails to avoid manipulation..

These patterns are consistent with recognized behavioral mechanisms such as those outlined in the MINDSPACE framework, where messengers, defaults, and salience shape how individuals interpret and engage with new practices.

The behavioral insights data described by interviewees and reflected in public commentary also showed that patterns such as reciprocity, connector centrality, and local hesitation could be identified through relational or behavioral signals. These patterns can be represented quantitatively and incorporated into dashboards that allow leaders to observe where influence, confidence, and friction reside. This information provides a basis for targeted interventions that reinforce trust and responsible use.

Together, these insights address the second research question by showing that AI adoption emerges from the interaction of influence structures and timely behavioral prompts, as in RQ2.1 and RQ2.2. Adoption accelerates when employees can observe credible examples, when peers they trust encourage experimentation, and when the organization creates clear and supportive cues and positive choice options within daily work. These conditions strengthen both AI literacy and responsible engagement across teams.

1.1.3 Learning Dynamics: AI Readiness as a Continuous Socio-Technical System

Across interviews and public commentary, a consistent insight emerged that AI readiness is not a fixed collection of employee skills but a dual capacity, shaped by individual adaptability and learning motivation as well as an organizational capability reinforced through continuous learning. Participants described how employees develop confidence when they can observe real outcomes, adjust their approaches, and receive ongoing guidance. In practice, readiness often depends on whether teams can identify breakdowns such as silos, unclear information, or blocked handoffs, and whether trusted peers and connectors can help restore clarity and coordination through local, workflow-embedded learning. These observations align with the scholars in the literature review's view that people-side knowledge management functions as an underlying infrastructure for adaptive learning and that relational and behavioral signals are essential for interpreting how knowledge actually moves across an organization.

1. **Continuously Reflecting on evidence**: Interviewees frequently characterized learning as iterative and grounded in evidence from daily work. Employees gained proficiency when they could experiment with AI within their own workflows, compare outputs with expectations, and refine their understanding through repeated use. Public commentary echoed this view, emphasizing that fluency develops through cycles of interaction and reflection rather than through one-time training sessions. Together, these perspectives suggest that readiness is strengthened through feedback loops that connect experimentation with evolving organizational context.
2. **Reinforcing Psychological Safety**: Participants also noted conditions that can disrupt these learning cycles. Unclear intent, inconsistent communication, and uncertainty about appropriate boundaries created hesitation during early exploration. Limited psychological safety further discouraged employees from testing new tools or sharing what they learned. These barriers reflect broader findings in the literature that adaptive learning requires shared meaning, visible norms, and explicit governance structures that help individuals understand how their actions contribute to organizational goals.
3. **Integrating Signals**: Several interviewees and public commentators pointed to the importance of leveraging organizational data to support this continuous learning. Modern organizations generate substantial telemetry, workflow data, and behavioral insights through digital tools, surveys, and social listening. When interpreted responsibly, these data can help leaders identify where learning is progressing, where employees encounter friction, and where additional support is needed. This aligns with the literature on people-side KM, which emphasizes the value of integrating both relational and behavioral signals into the ongoing management of knowledge systems. Within this view, approaches such as Organizational Network Analysis and behavioral analytics do not replace governance but offer a way to observe how learning unfolds in practice. They provide visibility into trust flows, collaboration patterns, informal expertise pathways, and emerging continuity risks, helping organizations recalibrate norms and processes as conditions shift. These insights reflect the broader socio-technical perspective in which knowledge is shaped not only by what is documented but also by how people interact, coordinate, and make sense of new tools.

The findings reinforce long-standing principles in socio-technical systems and distributed cognition, which emphasize that learning emerges from the interactions of people, information,

and tools. These theoretical perspectives help explain why organizations that align social context with technical systems demonstrate faster, more confident AI adoption.

Taken together, the findings show that AI readiness develops through the interaction of workflow context, social learning, and psychological safety, supported by organizational systems that make learning patterns visible. When organizations connect purposeful experimentation with real work and use data-driven diagnostics to understand how learning unfolds, people-side knowledge management becomes part of a continuous and self-reinforcing system for AI-driven transformation. In this sense, readiness is sustained through ongoing cycles of observation, interpretation, and adjustment that integrate governance clarity with everyday practice.

This pattern was further reflected in temporal coordination data from multiple deployments summarized in Appendix A. Across more than 1,500 organizations observed between 2020 and 2025, team-level pacing norms repeatedly appeared as measurable behavioral signatures shaping knowledge flow and execution speed. These observations provide the empirical foundation for the following insight.

- **Shared pacing norms:** A similar dynamic appeared in the temporal rhythms of teams and in how they absorbed new connections and newcomers. Shared pacing norms, including how quickly groups respond, how they sequence preparation, and how often they cycle through coordination, were not merely emergent patterns.
- **Reinforced by informal influencers:** They were reinforced by informal influencers whose behavior set expectations for the broader network.
- **Divergence creates latency:** When these rhythms diverged across teams, handoffs slowed and interpretations drifted, increasing knowledge latency even when formal workflows were well defined.
- **Practical probes:** In LEAD deployments, simple socialization and onboarding formats such as structured introductions, shadow calls, and short in-team training rotations served as practical probes of these dynamics.
- **Telemetry as signal:** When teams were able to pull newcomers and newly connected colleagues into real work quickly, these touchpoints triggered active exchange and visible learning. When inertia was high due to either apathy or lack of trust, the same prompts resulted in superficial contact, limited follow-through, or stalled participation. Across deployments, telemetry on participation and follow-up in these interactions provided a quantitative signal of team-level inertia and readiness. Taken together, these patterns reinforce that workflow efficiency and learning capacity are socio-behavioral as much as procedural, shaped by how teams collectively pace, integrate, and mobilize their members.

1.1.4 Translating Behavioral Evidence into Measurable Indicators

Both the literature review and the case evidence in Appendix A show that people-side KM can be represented not only through qualitative descriptions but also through measurable indicators that approximate how knowledge flows. Across more than 1,500 organizations, relational and behavioral patterns such as cross-functional responsiveness, trust-confirmed ties, and recurring interaction rhythms were translated into indicators like connector density, trust density, and knowledge latency. These indicators do not replace interpretation; rather, they offer a

standardized way to observe whether knowledge pathways are strong, fragile, or unevenly distributed.

As described in the literature review (Part 4.3) and elaborated conceptually in Appendix A, A3, this translation requires integrating:

- Passive ONA, which captures structural opportunities for interaction
- Active ONA, which captures relational qualities such as trust, perceived expertise, and psychological safety.

Trust density, for instance, bridges these domains by weighting structural ties with confirmed relational strength. Knowledge latency likewise differs from structural graph measures such as diameter or shortest-path distance; instead, it captures the observed time it takes for information, assistance, or decisions to move through real coordination rhythms. The behavioral model also incorporates concepts analogous to centrality, such as bridging influence and timely access to expertise, because these constructs are established predictors of diffusion, learning, and organizational adaptability.

Positioning these constructs as indicators has two implications for AI readiness.

- First, it enables organizations to monitor learning and adoption using behavioral evidence rather than relying solely on surveys or assumptions embedded in hierarchy.
- Second, it demonstrates how people-side KM can generate machine-usable signals that AI systems can incorporate as features, for example to prioritize routing, tailor nudges, or detect emerging bottlenecks, while maintaining privacy through aggregation.

In this sense, the indicators reviewed in the literature and synthesized in Appendix A function as an operational layer connecting human behavior, knowledge flow, and AI-supported decision making. Together, these insights point to the need for an integrated framework capable of connecting human behavior, information flow, and workflow signals into a single continuous learning system. These measurable indicators provide the conceptual and empirical foundation for the framework introduced in the next section, which articulates how such a system can operate in practice.

2. The Organizational Intelligence Loop (OIL) Framework

The combined dataset and interviews revealed four recurring gaps that limit sustainable AI adoption:

- Knowledge lacks traceability or ownership (**information gap**).
- Informal relational context remains invisible (**relational gap**).
- Feedback between workflow and learning is missing (**process gap**).
- AI systems lack two layers required for adaptation: (1) organizational context and guardrails, such as expertise maps, access boundaries, authority signals, and real

coordination pathways; and (2) intervention pathways, such as nudges, workflow adjustments, and governance actions that translate detected friction into behavioral change. (**adaptation gap**).

To address these interconnected gaps, this study introduces the Organizational Intelligence Loop (OIL), a socio-technical model that links People, Information, and Process layers through an adaptive Agentic AI Design layer to form a continuous learning system. The framework is grounded in the empirical observation that trust, influence, and learning operate as one interdependent system rather than as separate initiatives.

Research Question	OIL	Purpose in the System
Behavioral Insights (H2) How do behavioral and network data reveal influencers, bottlenecks, relational biases, and trust patterns, and how can targeted interventions strengthen literacy and confidence?	**People Layer**	Makes individual responses (fear, uncertainty, confidence, identity threat, perceived safety) and relational dynamics (trust flows, influence patterns, bottlenecks) visible; enables targeted interventions that strengthen psychological safety and support responsible AI adoption.
Curation & Governance (H1) How can information architecture, metadata standards, and governance practices be designed and socialized to support transparency, traceability, and trust?	**Information Layer**	Creates a unified, validated, and accountable knowledge base; ensures metadata is current, traceable, and socially verified; foundation for reliable human and AI reasoning.
Continuous Learning Systems (H3) How can people-side KM insights be translated into continuous learning systems that support readiness for evolving AI tools?	**Process Layer**	Embeds experimentation, reflection, and adaptive feedback into workflows; connects human learning patterns with AI reasoning to form a continuous system of improvement.
Derived from H1-H3	**Agentic AI Design Layer**	Integrates feedback across people, information, and workflows; adapts reasoning and outputs based on behavioral and process signals; forms the continuous machine-learning engine that reinforces organizational learning.

This mapping grounds OIL in empirical observation: trust, influence, and learning operate as one interdependent system rather than as separate initiatives. Several barriers identified in the interviews lie beyond the direct reach of any framework, including lack of leadership intent or clarity of purpose. For example, as noted in 1.1.1, interviewees described cases where executives may launch AI initiatives without clear problem definition or shared success metrics. Internal politics and power asymmetries, such as misaligned incentives between IT and business units, can also block collaboration despite technical readiness.

Although OIL cannot directly solve these leadership and governance failures, it can expose and contextualize them by making patterns of expertise, trust, and influence visible. Even cultural resistance and identity threats, which are deeply human, can be partially mapped through the

People Layer. This allows organizations to detect where bias, mistrust, fear, (mis)influence, or misinformation accumulate and to design interventions that support transparency and psychological safety.

Each research question aligns with a structural layer of the OIL framework. The People Layer integrates both individual behavioral responses (fear, uncertainty, confidence, psychological safety) and social dynamics (trust flows, influence networks, relational bottlenecks), reflecting interview evidence interviewees' and commentators' insight that AI adoption is shaped by how people think, feel, and interact, not only by formal structures.

2.1 Layer Definitions

The Organizational Intelligence Loop (OIL) frames enterprise AI as a continuous feedback system connecting People, Information, and Process layers through an adaptive Agentic AI design layer.

Each layer addresses a key question:

1. **People Layer:** Who knows what, behaves how, and who influences whom?
2. **Information Layer:** What is known, and how traceable and trustworthy is it?
3. **Process Layer:** How does work, knowledge, information and decision-making move through the organization?
4. **Agentic AI Design Layer:** How does it adapt and reinforce learning?

Together, they form a closed feedback loop aligning social networks, knowledge systems, and AI adaptation into one framework for collective intelligence. OIL embodies the insight that sustainable AI transformation depends on parallel learning between humans and machines, where each layer informs and strengthens the others.

- **Illustration 1: IS view**

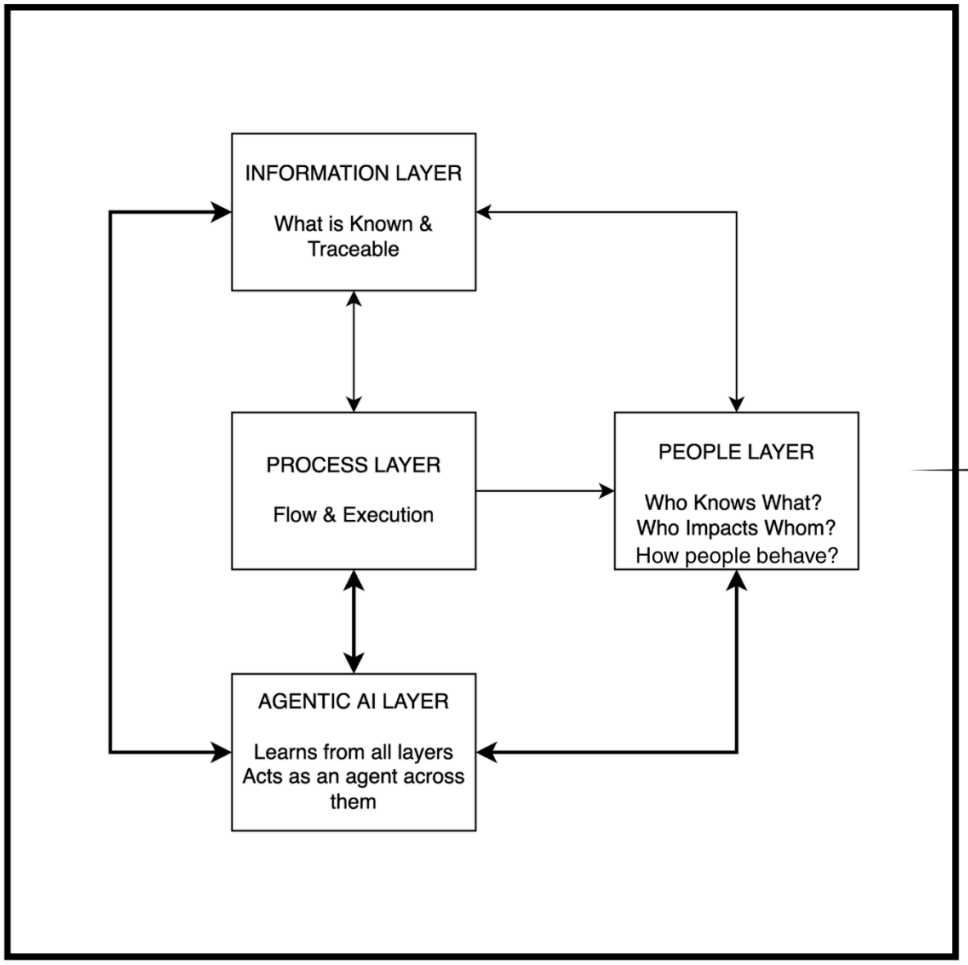

High-performing, high-fit organization

2.2 People Layer — Who Knows What, How People Behave, Who Influences Whom?

The People Layer captures the human dynamics through which knowledge is created, interpreted, and applied. OIL treats these dynamics not as informal background activity but as a measurable system grounded in behavioral evidence, influence patterns, and trust relationships. This layer clarifies who truly enables learning—not only who has expertise, but whose actions shape interpretation, confidence, and coordination across teams.

Drawing on findings from Parts 2 and 3 and from large-scale observations in the LEAD dataset, the People Layer brings together two domains of insight:

• **Cognitive and behavioral responses:** fear, uncertainty, confidence, psychological safety, willingness to experiment (bias, norms…)
• **Social and relational structures:** influence networks, trust flows, bottlenecks, hidden brokers, cross-boundary connectors

These factors determine whether knowledge moves through the organization, whether AI outputs are understood and trusted, and whether employees feel safe enough to explore new tools.

Existing enterprise platforms such as Microsoft Viva Insights reveal collaboration traces such as meetings, shared channels, and responsiveness rhythms, but these traces rarely function as a coherent learning system. Within OIL, passive ONA (structural traces) is interpreted alongside active ONA (trust-confirmed ties, perceived expertise, reciprocity, and informal influence). Importantly, OIL does not treat ONA as a stand in for all behavioral insights. Where mechanisms such as bias, norms, messenger effects, or the impact of embedded prompts and defaults are not inferable from network structure, the People Layer can incorporate lightweight active inputs (for example short pulse questions) and adoption or workflow telemetry that makes hesitation, reliance patterns, and behavior change observable and testable.

To make these signals usable by AI systems, the People Layer can be represented as a typed organizational graph that links people, roles, teams, and work artifacts, with edges populated by passive ONA, active ONA, and bounded telemetry. A graph-based retrieval layer can then provide agentic systems with a policy-aware slice of People context at runtime, meaning during task execution when the agent is generating an answer or taking an action (not during model training), so recommendations are grounded in who knows what, who influences whom, and where coordination stalls, rather than assuming this context is native to the model: most foundation models are trained on general data and do not contain an organization’s current expertise map, trust ties, decision rights, or access boundaries unless those signals are explicitly instrumented and retrieved from enterprise systems under governance controls.

This visibility makes it possible to identify connectors and brokers who accelerate learning, trust hubs that legitimize new practices, structural gaps that slow diffusion, and hidden contributors

whose departure would jeopardize tacit knowledge flow. Interviewees repeatedly noted that employees look to trusted peers rather than formal leaders when deciding whether AI use is safe, useful, or appropriate for their work. The People Layer makes these informal pathways visible. It supports practical actions such as pairing credible early adopters with hesitant groups, involving trusted messengers in pilots, and highlighting peer examples that demonstrate value.

The People Layer ultimately functions as the social engine of the OIL framework. It translates tacit human behavior into observable signals of trust, influence, and learning capacity. When these conditions are strong, the organization operates as a high-fit system in which the right relationships and norms support continuous adaptation. By grounding human dynamics in structured evidence, the People Layer provides the connective tissue that allows the Information Layer and the Process Layer to operate as an integrated system of organizational intelligence.

2.3 Information Layer — What Is Known, and How Traceable and Trustworthy Is It?

The Information Layer represents the organization's formal knowledge base: the data, documents, and structured records stored across systems such as Salesforce, Zendesk, Confluence, and internal repositories. These platforms form the organization's informational memory, but they do not indicate whether content is accurate, current, or trusted. Most enterprise AI systems therefore lack the ability to evaluate credibility or relevance without additional context. In OIL, "information" also includes curated organizational signals, including bounded People-layer outputs such as network structure, expertise signals, and trust-confirmed influence pathways, because these signals shape retrieval and downstream agent behavior and therefore must be governed as first-class information assets.

Within the OIL framework, the Information Layer is strengthened through dynamic metadata governance that links content and organizational signals to the social and epistemic evidence captured in the People Layer. OIL traces who creates, edits, or curates information; whose expertise or influence lends weight to certain assets; and where trust, recency, endorsement, or verified stewardship contribute to the perceived reliability of knowledge. Content authored or maintained by credible contributors gains evidentiary weight, while outdated, duplicated, or low-engagement material becomes visible as a potential risk to organizational coherence.

These provenance and trust signals refresh continuously as employees engage with information, consulting, correcting, or validating it in the course of their work. Over time, the Information Layer accumulates a living history of maintenance, authorship, credibility, and consensus. This shifting metadata landscape allows the system to surface authoritative sources, highlight inconsistencies, and detect early signs of informational fragmentation when teams diverge in how they maintain core knowledge.

By grounding formal content and curated organizational signals in these evolving social and epistemic signals, the Information Layer transitions from a static repository into a traceable and trustworthy knowledge graph. It enables AI to reason not only over content, but also over the organization's confidence in that content and the bounded context needed for responsible, policy-aware intelligence.

[2]2.4 Process Layer — How Does Work, Knowledge, Information, and Decision-Making Move Through the Organization?

The Process Layer describes the pathways through which work, knowledge, information, and decisions move across teams. It connects intent to execution by making coordination mechanisms visible, including handoffs, approvals, escalation routes, and points of accountability. Enterprise systems such as ServiceNow, Jira, and Workday already capture parts of this activity through tickets, tasks, and status changes, but they rarely show whether coordination is actually functioning well, who is driving decisions, or where behavior and interpretation break the workflow.

This layer connects to the study's findings in two ways. First, interviewees described adoption breakdowns as coordination failures, not just tool limitations: unclear intent, inconsistent guardrails, and siloed decision rights created hesitation and slowed experimentation. Second, many of the RQ1 governance problems surfaced as failures of workflow ownership. Employees often could not tell who was accountable for updating knowledge, approving use cases, or resolving cross-system conflicts. In practice, trust depended on whether these pathways were visible and consistently enforced, so employees could understand how decisions were made and how exceptions were handled.

Within OIL, the Process Layer operationalizes governance by turning policies and knowledge rules into routinized actions. When process telemetry is interpreted alongside the People and Information layers, the organization can detect where approvals stall, where handoffs repeatedly fail, where teams diverge in how they interpret the same information, and where decisions concentrate in overloaded individuals. This visibility supports targeted interventions such as clarifying decision rights, redesigning workflows, and embedding low-friction guidance in the flow of work, aligned with the specific failure modes surfaced in the findings.

By linking process signals with social and informational context, the Process Layer transforms workflow data into a diagnostic view of organizational trust and execution. It helps explain why certain workflows succeed or stall, how decisions propagate, and how knowledge is applied as it moves between actors. Over time, it functions as a continuous coordination and learning system that allows organizations to adapt routines as conditions change while maintaining coherence across people, information, and work.

2.5 Agentic AI Design Layer: How does it adapt and reinforce learning?

Most enterprise AI systems rely almost entirely on the Information Layer. They are trained on documents, wikis, structured databases, and other static content stored in platforms such as Confluence, Notion, or Salesforce. In this configuration, AI is used mainly for retrieval, summarization, and drafting. Although employees may refine prompts or offer corrections, this feedback operates at the surface and does not influence the system's deeper understanding of how knowledge is created, shared, or interpreted inside the organization.

This separation produces a familiar pattern in many programs. When employees doubt the accuracy or relevance of AI outputs, they limit their use. Limited use prevents the system from

[2]

learning, which leads to increasingly generic or misaligned responses. Without behavioral or social feedback loops explicitly designed in, the AI remains disconnected from how the organization actually works and fails to improve over time.

The Agentic AI Design Layer in the OIL framework addresses this limitation by linking AI systems to continuous signals from across the enterprise. At runtime, the agent retrieves bounded, policy-aware context from each layer rather than assuming it is native to the model. Instead of learning only from documents, the model receives structured insight from the People Layer, where patterns of expertise, influence, trust, and adoption are visible. It also receives provenance and metadata freshness from the Information Layer and real-time coordination and workflow signals from the Process Layer. Through secure and privacy-aware connectors, the AI system gains visibility into how knowledge moves, who maintains it, how employees behave, how decisions are made, and where processes succeed or encounter friction. Where permitted, the agent can also trigger governed write-backs, such as routing a question to an identified expert, flagging a knowledge asset for review, or initiating a workflow step.

With this expanded context, Agentic AI becomes a co-learning system rather than a static tool. It can recognize where trust is strong or weak, identify emerging knowledge gaps, detect fragmentation across teams, and adjust its recommendations accordingly. It can highlight when content requires curation, when experts should be connected, or when workflows show signs of overload. These adaptations reflect not only patterns in documents but the social and operational conditions in which knowledge is used, made explicit in the other layers.

Conceptually, OIL reframes enterprise AI from a model that is fine-tuned on fixed datasets to one that learns continuously from the organization’s lived activity. As employees interact with AI, the system improves its contextual accuracy. As accuracy increases, employees engage more often, producing richer behavioral signals that further enhance the model’s relevance. This creates a reinforcing learning cycle in which human and machine adaptation evolve together. Over time, the Agentic AI Design Layer becomes an engine of organizational learning. It strengthens transparency, supports responsible use, and amplifies the collective intelligence of the enterprise by ensuring that AI is grounded in how people work, how they behave, how knowledge is maintained, and how decisions are made.

For the OIL framework to operate beyond a conceptual model, its layers must be supported by a coherent technical structure that enables continuous learning across people, information, and workflows. The following section outlines the system architecture and theoretical foundations that make this possible, translating the framework’s principles into mechanisms that can support real-time organizational intelligence.

2.6 System and Foundations Informing the Architecture

The OIL builds on the findings and theoretical perspectives outlined earlier by proposing a multi-layered architecture that links human behavior, information systems, and workflow activity through a unified semantic and technical design. The purpose of this structure is to make organizational learning continuous, transparent, and measurable. As the literature suggests, effective learning systems emerge not from isolated tools but from the coordinated interaction of

people, information, and processes. The architecture below represents an operational translation of these principles.

2.2.1 System Architecture

OIL is implemented through interoperable components that allow people-side knowledge, informational assets, and process telemetry to reinforce one another. Rather than introducing a new platform, the architecture provides a shared foundation through which existing enterprise systems can exchange context and meaning. Three design elements support this integration.

Unified identity graph.
This component connects expertise, collaboration patterns, and trust relationships into a coherent representation of the organization. By drawing on behavioral and relational evidence outlined in earlier sections, including lightweight active inputs such as pulse surveys or one question prompts when affect or trust must be explicitly captured, it creates the semantic backbone through which knowledge and decision flows can be contextualized under role-based access and privacy safeguards.

Append-only event streams.
These streams capture real-time workflow and communication patterns from systems such as Jira, Salesforce, and ServiceNow. They provide temporal signals about how work progresses, how information moves, and where decisions are enacted. In doing so, they reflect the learning dynamics highlighted in the findings, where behavior and context evolve continuously rather than through discrete interventions.

Shared metadata schema.
This schema standardizes how knowledge assets are described, including their source, ownership, provenance, freshness, and trust. It transforms traditional repositories into living knowledge ecosystems that mirror how employees create, update, and interpret information. This directly supports the governance and curation challenges raised in earlier research questions.

Together, these mechanisms allow the four layers of OIL to function as an integrated learning system.

- The People Layer captures behavioral and influence patterns.
- The Information Layer maintains knowledge that is current, traceable, and socially validated.
- The Process Layer embeds workflow context and signals of coordination.
- The Agentic AI Design Layer synthesizes these insights into adaptive, explainable outputs by retrieving bounded context from the other layers at runtime. Through shared feedback across these layers, organizations develop a self-improving loop in which both humans and systems evolve together.

2.2.2 Foundations Informing the Architecture

The architecture draws on a range of theoretical traditions introduced in the literature review. Socio-technical systems theory emphasizes that organizational effectiveness depends on the alignment of human and technical subsystems (Trist & Emery, 1951). Distributed cognition highlights how knowledge resides across people, tools, and representations rather than within individuals alone (Hutchins, 1995). Learning-organization research demonstrates that sustained improvement requires ongoing feedback, shared meaning, and adaptive routines (Argyris & Schön, 1978; Nonaka & Takeuchi, 1995).

Network science informs the relational dimension of OIL by showing how influence, connectivity, and structural positions determine whether knowledge diffuses or becomes blocked (Freeman, 1978; Borgatti et al., 2018). Behavioral frameworks such as MINDSPACE and ADKAR contribute insight into how norms, incentives, and perceived credibility shape adoption and readiness.

The architecture also incorporates principles from semantic and graph-based knowledge management. Traditional semantic models such as RDF, OWL, and SKOS offer formal structures for classifying entities and relationships, yet they remain static and do not account for the behavioral context in which knowledge is interpreted. OIL preserves the precision of these models while extending them with live signals from collaboration and workflow, thereby bridging the gap between classification and real-world use.

2.2.3 Standards Alignment and Broader Application

The OIL architecture aligns with ISO 30401 by operationalizing principles of knowledge ownership, traceability, and currency. It reflects the NIST AI Risk Management Framework's emphasis on context, accountability, and explainability by ensuring that AI systems are grounded in transparent human and organizational signals. Because OIL incorporates people side evidence, it also aligns with enterprise security and privacy expectations by supporting controls commonly required under ISO IEC 27001 and ISO IEC 27701, and audit oriented assurance frameworks such as SOC 2. In practice, this means designing for least privilege access, purpose limitation, data minimization, retention controls, and auditable provenance so that behavioral signals can be used without expanding exposure of sensitive personal data. Through measurable constructs such as trust, influence, expertise, and collaboration patterns, OIL establishes a shared language across behavioral analytics, KM systems, and AI models.

By integrating structured knowledge models with real-time behavioral data, as well as publicly recognized standards, OIL reframes knowledge management as a living organizational intelligence system. It provides a foundation for enterprise-wide generative-AI adoption by aligning human behavior, knowledge flow, and computational learning within a single architecture.

Notably, beyond AI, the same principles support large-scale change management, organizational redesign, and cultural transformation, offering leaders a dynamic evidence base for understanding how their organizations learn, adapt, and evolve over time. We leave this to be pursued in further research.

Limitations and Future Research Directions

1. Agentic AI Governance and Alignment Systems

While the Organizational Intelligence Loop (OIL) enhances organizational learning by making trust relationships, expertise signals, and workflow context observable, several challenges lie beyond the scope of what OIL alone can resolve. These challenges highlight important directions for future research.

A central issue is the governance of increasingly autonomous AI systems. Current generative and agentic models typically operate without persistent organizational context by default: they do not inherently know which sources are authoritative, who holds the right access, or whose domain knowledge is most reliable (National Institute of Standards and Technology [NIST], 2023). As industry leaders such as Box CEO Aaron Levie have noted, this absence of contextual grounding introduces alignment instability, authorization, governance and security vulnerabilities. Even with fine-tuning or human-feedback loops, misalignment cannot be fully solved at the model level, since human inputs themselves may reflect conflicting incentives or unrecognized biases(NIST, 2023; Shivpuja, 2025). Recent studies suggest that these tensions can produce behavioral drift (including instruction or alignment drift), where an agent's outputs gradually diverge from intended organizational constraints and values over time, especially across extended interactions or subsequent fine-tuning (Clymer, Westover, & Khandelwal, 2025; Dongre et al., 2025).[3]

OIL helps surface the social and informational conditions that shape trustworthy collaboration, but it does not determine how real-time oversight should be distributed between humans and AI. Organizations must still evaluate when humans should act as supervisors of system behavior, when they should serve as contextual anchors, and how these roles evolve as AI systems gain autonomy. Both approaches carry trade-offs. Continuous human review protects accountability but introduces cognitive burden and inconsistency, while reducing humans to contextual providers allows scale but risks disconnecting AI actions from evolving organizational intent (NIST, 2023; Shivpuja, 2025).

These gaps point toward a broader research agenda on multi-agent governance. Scholars such as Tsamados and colleagues argue that future AI systems will require distributed oversight structures in which human and machine agents collaborate to maintain alignment, while recent work on AI personality drift underscores the need for continuous monitoring of agent behavior over time (Tsamados et al., 2024; Siqueira de Cerqueira et al., 2025). One promising direction is a multi-agent extension of OIL, in which AI agents critique, verify, and audit one another's reasoning while humans supply the normative, contextual, and ethical grounding that anchors system behavior. Such architectures would move human–AI governance from a hierarchical model to a reciprocal one, where mutual accountability replaces unilateral supervision (NIST, 2023).

3

2. Data, Measurement, and Retrieval Operationalization Limits

This study defines OIL as a socio-technical architecture that connects People, Information, and Process layers through an Agentic AI layer and describes how these layers exchange signals through feedback loops. It also outlines enabling components such as a unified organizational identity graph, append-only event streams, and shared metadata structures, positioning semantic and metadata work as part of the governance foundation that OIL extends with behavioral and workflow context.

However, the study remains at an architectural level and does not implement or empirically validate the end-to-end representation and retrieval stack required to operationalize OIL inside production enterprise agentic workflows. In particular, it does not provide

- (1) a production-grade specification of the organizational graph (for example node and edge types, identity resolution, temporal updates, and scaling constraints),
- (2) the retrieval and grounding procedures needed for agents to use this context consistently and safely (for example query, ranking, provenance handling, and conflict resolution across layers),
- or (3) the deployment governance and evaluation requirements (for example access controls, auditability, monitoring, and robustness testing).

Future research should translate OIL from an architectural model into a deployable, policy-aware context representation and retrieval layer and evaluate its effects on agent reliability, safety, and adoption outcomes under real organizational constraints.

3. Organizational Diagnostics for AI Adoption Outcomes and Performance Translation

This study does not attempt to link behavioral indicators directly to business performance metrics (e.g., revenue, customer impact, operational efficiency). This reflects both the methodological constraints of enterprise research and the conceptual role of socio-technical diagnostics: behavioral and network indicators function as an organizational “imaging system” that reveals hidden dependencies, bottlenecks, and readiness conditions, but does not itself determine outcomes.

At the same time, the study does demonstrate that behavioral evidence can be operationalized into a trainable model that produces consistent, machine-usable indicators (e.g., trust density, connector density, knowledge latency) and can support predictive or monitoring use cases at the level of adoption dynamics, coordination friction, and learning conditions. In this sense, the contribution is not merely descriptive: it shows that socio-technical states can be modeled in ways that are repeatable, comparable across units, and usable for intervention design.

Two organizations facing similar diagnostic patterns may choose different interventions, operate under different constraints, or differ in leadership alignment and cultural flexibility. As a result, performance improvements should not be assumed to follow uniformly from diagnostic visibility alone.

Study-level boundary conditions and data limits

The empirical evidence in this study is derived from organizations that generate behavioral and collaboration signals through LEAD deployments; this may bias coverage toward enterprise environments already willing to instrument knowledge-sharing behavior. In addition, the dataset does not support industry-specific or size-specific parameterization, and observed patterns may not generalize to organizations with materially different collaboration infrastructure, governance maturity, or regulatory constraints.

Future research should pair diagnostic signals with longitudinal intervention designs and carefully aggregated, privacy-preserving outcome measures to clarify when (and under what conditions) diagnostic insights translate into measurable operational impact.

4. Behavioral Nudges, Autonomy, and Ethical Trade-offs

This study treats nudges as practical, evidence-based mechanisms for reducing friction and supporting responsible AI adoption in the flow of work. However, this study does not evaluate how nudges affect perceived autonomy in enterprise settings, including when choice architecture feels supportive versus coercive. Because this study focuses on socio-technical readiness and adoption dynamics, it does not evaluate the ethical boundary conditions of nudges—such as when interventions become manipulative, how consent and transparency should be handled, or how power asymmetries (e.g., manager–employee dynamics) shape perceived voluntariness. At the same time, organizations routinely shape behavior through policies, performance feedback, managerial conversations, training, and incentives; in this sense, nudges are best understood as a more granular, workflow-embedded extension of existing change mechanisms.

Future research should therefore develop and test governance principles for autonomy-preserving nudges in enterprise AI contexts, including disclosure standards, meaningful opt-out and contestability mechanisms, proportionality to risk, and auditability safeguards to distinguish supportive guidance from coercive behavioral control and to assess effects on perceived autonomy, trust, and adoption outcomes.

Conclusion

This study shows that effective AI adoption is not primarily a technical challenge but an organizational learning challenge. Across interviews, public commentary, and large-scale evidence from LEAD deployments, the decisive variables were patterns of trust, human bias and autonomic behavior, informal influence, social diffusion, and psychological safety. When these human dynamics are visible and measurable, organizations gain the ability to manage transformation proactively rather than reactively.

The OIL synthesizes these insights into a framework that links behavioral, informational, and process layers into a continuous feedback system. By transforming collaboration patterns, network signals, and workflow activity into structured intelligence, OIL operationalizes the long-standing vision of knowledge management as a living system—one that adapts as people learn, as conditions change, and as new technologies emerge.

These findings reinforce theoretical perspectives from socio-technical systems, behavioral insights, distributed cognition, and organizational learning, which all emphasize that individual learning arises from the interaction of people, processes, and information – and that the organization, itself, can learn. OIL extends these traditions by offering a practical architecture through which organizations can monitor learning dynamics, strengthen knowledge flows, and sustain readiness over time.

As enterprises move toward more agentic forms of AI, the core requirement remains the same: alignment between human context and system design and behavior. We argue further that any multi-agent AI ecosystem will depend on the same foundations that support effective human networks—shared context, mutual trust, and mechanisms for continuous correction. Integrating OIL's feedback logic into AI governance and coordination provides a pathway for ensuring that intelligent agents remain transparent, accountable, and aligned with organizational values.

In essence, AI adoption becomes knowledge management at scale. When organizations treat learning as a continuous socio-technical process, people, processes, and intelligent systems evolve together. OIL offers a foundation for this evolution, enabling enterprises to build the adaptive capacity required for responsible, resilient, and long-term AI transformation.

Reference

American Institute of Certified Public Accountants. (2017). SOC 2®—Reporting on controls at a service organization relevant to security, availability, processing integrity, confidentiality, or privacy. AICPA.

Andreessen Horowitz. (2023). The dynamics of network effects. https://a16z.com/the-dynamics-of-network-effects/

Argyris, C., & Schön, D. A. (1978). Organizational learning: A theory of action perspective. Addison-Wesley. https://www.scirp.org/reference/referencespapers?referenceid=697693

Ashford, S. J., & Black, J. S. (1996). Proactivity during organizational entry: The role of desire for control. Journal of Applied Psychology, 81(2), 199–214. https://doi.org/10.1037/0021-9010.81.2.199

Baumard, P. (1999). Tacit knowledge in organizations. SAGE.

Bauer, T. N., & Erdogan, B. (2011). Organizational socialization: The effective onboarding of new employees. In S. Zedeck (Ed.), APA handbook of industrial and organizational psychology (Vol. 3, pp. 51–64). American Psychological Association. https://psycnet.apa.org/record/2010-06019-002

Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. Interacting with Computers, 23(1), 4–17.

Biemer, P. P., Chen, P., & Wang, K. (2013). Using level-of-effort paradata in non-response adjustments with application to field surveys. Journal of the Royal Statistical Society: Series A, 176(1), 147–168. https://academic.oup.com/iwc/article-abstract/23/1/4/693091?redirectedFrom=fulltext

Borgatti, S. P., Everett, M. G., & Johnson, J. C. (2018). Analyzing social networks (2nd ed.). SAGE Publications.

Bosnjak, M., & Tuten, T. L. (2003). Prepaid and promised incentives in web surveys: An experiment. Social Science Computer Review, 21(2), 233–252. https://journals.sagepub.com/doi/10.1177/0894439303021002006

Burt, R. S. (1992). Structural holes: The social structure of competition. Harvard University Press.

Cable, D. M., Gino, F., & Staats, B. R. (2013). Breaking them in or eliciting their best? Reframing socialization around newcomers' authentic self-expression. Administrative Science Quarterly, 58(1), 1–36. https://doi.org/10.1177/0001839213477098

Cialdini, R. B. (2007). Influence: The psychology of persuasion (Rev. ed.). Harper Business.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. Annual Review of Psychology, 55, 591–621. https://www.researchgate.net/publication/8902776_Social_Influence_Compliance_and_Conformity

Contractor, N. S., Wasserman, S., & Faust, K. (2006). Testing multitheoretical, multilevel hypotheses about organizational networks: An analytic framework and https://journals.aom.org/doi/10.5465/amr.2006.21318925

Cooper-Thomas, H. D., & Anderson, N. (2006). Organizational socialization: A new theoretical model and recommendations for future research and HRM practices in organizations. Journal of Managerial Psychology, 21(5), 492–516. https://www.researchgate.net/publication/228345225_Organizational_Socialization_A_New_Theoretical_Model_and_Recommendations_for_Future_Research_and_HRM_Practices_in_Organizations

Corbett, A. T., Koedinger, K. R., & Anderson, J. R. (1997). Intelligent tutoring systems and adaptive learning. User Modeling and User-Adapted Interaction, 7(3), 253–278. https://act-r.psy.cmu.edu/wordpress/wp-content/uploads/2012/12/173Chapter_37_Intelligent_Tutoring_Systems.pdf

Clymer, J., Westover, A., & Khandelwal, A. (2025). *Will AI systems drift into misalignment?* AI Alignment Forum. https://www.alignmentforum.org/posts/u8TYRhGPD878i3qkc/will-ai-systems-drift-into-misalignment

Cross, R., Borgatti, S. P., & Parker, A. (2002). Making invisible work visible: Using social network analysis to support strategic collaboration. California Management Review, 44(2), 25–46. https://www.researchgate.net/publication/240104637_Making_Invisible_Work_Visible_Using_Social_Network_Analysis_to_Support_Strategic_Collaboration

Cross, R., & Parker, A. (2004). The hidden power of social networks: Understanding how work really gets done in organizations. Harvard Business School Press.

Cross, R., Parker, A., Prusak, L., & Borgatti, S. (2001). Knowing what we know: Supporting knowledge creation and sharing in social networks. Organizational Dynamics, 30(2), 100–120. https://scholars.uky.edu/es/publications/knowing-what-we-know-supporting-knowledge-creation-and-sharing-in/

Cross, R., Rebele, R., & Grant, A. (2016). Collaborative overload. Harvard Business Review, 94(1–2), 74–79.https://hbr.org/2016/01/collaborative-overload

Cross, R., & Thomas, R. (2011). Driving results through social networks: How top organizations leverage networks for performance and growth. Jossey-Bass.

Cross, R., & Thomas, R. J. (2011, July). A smarter way to network. Harvard Business Review. https://hbr.org/2011/07/managing-yourself-a-smarter-way-to-network

Cummings, J. N., & Pletcher, C. (2011). Why project networks beat project teams. MIT Sloan Management Review, 52(3), 51–58. https://sloanreview.mit.edu/article/why-project-networks-beat-project-teams/

Dalkir, K. (2011). Knowledge management in theory and practice (2nd ed.). MIT Press.

Dalkir, K. (2017). Knowledge management in theory and practice (3rd ed.). MIT Press.

Davenport, T. H., & Prusak, L. (1998). Working knowledge: How organizations manage what they know. Harvard Business School Press.

Dolan, P., Hallsworth, M., Halpern, D., King, D., Metcalfe, R., & Vlaev, I. (2012). Influencing behaviour: The MINDSPACE way. Journal of Economic Psychology, 33(1), 264–277. https://www.sciencedirect.com/science/article/abs/pii/S0167487011001668

Dongre, V., Rossi, R. A., Lai, V. D., Yoon, D. S., Hakkani-Tür, D., & Bui, T. (2025). *Drift no more? Context equilibria in multi-turn LLM interactions* (arXiv:2510.07777). arXiv. https://arxiv.org/abs/2510.07777

Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. Administrative Science Quarterly, 44(2), 350–383. https://web.mit.edu/curhan/www/docs/Articles/15341_Readings/Group_Performance/Edmondson%20Psychological%20safety.pdf

Edmondson, A. C. (2019). The fearless organization: Creating psychological safety in the workplace for learning, innovation, and growth. Wiley.

Eraut, M. (2004). Informal learning in the workplace. Studies in Continuing Education, 26(2), 247–273. https://www.tandfonline.com/doi/full/10.1080/158037042000225245

Freeman, L. C. (1978). Centrality in social networks: Conceptual clarification. Social Networks, 1(3), 215–239. https://www.sciencedirect.com/science/article/abs/pii/0378873378900217

Gamba, D., Yu, Y., Yuan, Y., Schoenebeck, G., & Romero, D. M. (2024). Exit ripple effects: Understanding the disruption of socialization networks following employee departures. In

Proceedings of the ACM Web Conference 2024 (WWW '24). ACM. https://arxiv.org/abs/2402.15683

Gartner. (2025, November 19). Gartner identifies critical GenAI blind spots that CIOs must urgently address [Press release]. https://www.gartner.com/en/newsroom/press-releases/2025-11-19-gartner-identifies-critical-genai-blind-spots-that-cios-must-urgently-address0

Hackman, J. R. (2002). Leading teams: Setting the stage for great performances. Harvard Business School Press.

Hallsworth, M., & Kirkman, E. (2020). Behavioral insights. Behavioural Insights Team.https://www.bi.team/publications/behavioural-insights

Hansen, M. T. (2009). Collaboration: How leaders avoid the traps, build common ground, and reap big results. Harvard Business Press.

Hutchins, E. (1995). Cognition in the wild. MIT Press.

International Organization for Standardization. (2013). ISO/IEC 27001:2013 Information technology — Security techniques — Information security management systems — Requirements. ISO. https://pecb.com/en/whitepaper/iso-27001-information-technology-security-techniques-information-security-management-systems-requirements

International Organization for Standardization. (2018). ISO 30401:2018 knowledge management systems – Requirements. https://www.iso.org/standard/68683.html

International Organization for Standardization. (2019). ISO/IEC 27701:2019 Security techniques — Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management — Requirements and guidelines. ISO. https://www.iso.org/standard/71670.html

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. Academy of Management Annals, 14(1), 366–410. https://journals.aom.org/doi/10.5465/annals.2018.0174

Kosmyna, N., et al. (2025). Trust calibration in human–AI surveillance teams supported by a psychometric framework. TPM – Testing, Psychometrics, Methodology in Applied Psychology, 32(S2), 1847–1851. https://tpmap.org/submission/index.php/tpm/article/view/987

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task [Preprint]. arXiv. https://arxiv.org/abs/2506.08872

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

Lambe, P., & Milton, N. (2019). The knowledge manager's handbook: A step-by-step guide to embedding effective knowledge management in your organization (2nd ed.). Kogan Page.

McKinsey & Company. (2024). Embrace gen AI with eyes wide open. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/embrace-gen-ai-with-eyes-wide-open

McKinsey & Company. (2024). The state of AI in 2024. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024

McKinsey & Company. (2025). The learning organization: How to accelerate AI adoption. https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/the-learning-organization-how-to-accelerate-ai-adoption

Menlo Ventures. (2024). 2024: The state of generative AI in the enterprise. https://menlovc.com/2024-the-state-of-generative-ai-in-the-enterprise

Monge, P. R., & Contractor, N. S. (2003). Theories of communication networks. Oxford University Press.

Morrison, E. W. (2002). Newcomers' relationships: The role of social network ties during socialization. Academy of Management Journal, 45(6), 1149–1160.

National Institute of Standards and Technology. (2023). AI risk management framework. https://www.nist.gov/itl/ai-risk-management-framework

Nonaka, I., & Takeuchi, H. (1995). The knowledge-creating company: How Japanese companies create the dynamics of innovation. Oxford University Press.

Orlikowski, W. J., & Scott, S. V. (2014). What happens when evaluation goes online? Exploring apparatuses of valuation in the travel sector. Organization Science, 25(3), 868–891.

Phaladi, M., & Ngulube, P. (2022). Mitigating risks of tacit knowledge loss in state-owned enterprises in South Africa through knowledge management practices. South African Journal of Information Management, 24(1), a1462.

Prosci. (2006). ADKAR: A model for change in business, government and our community. Prosci Learning Center Publications.

Pugh, K. B. (2023). IKNS 5305 Module 2: Social network analysis vocabulary [Unpublished course handout]. Columbia University, School of Professional Studies.

Rogers, E. M. (2003). Diffusion of innovations (5th ed.). Free Press.

Ross, J. W., Weill, P., & Robertson, D. C. (2006). Enterprise architecture as strategy: Creating a foundation for business execution. Harvard Business School Press.

Schein, E. H. (2010). Organizational culture and leadership (4th ed.). Jossey-Bass.

Shao, Y., Zope, H., Jiang, Y., Pei, J., Nguyen, D., Brynjolfsson, E., & Yang, D. (2024). Future of work with AI agents: Auditing automation and augmentation potential across the U.S. workforce [Preprint]. arXiv. https://arxiv.org/pdf/2506.06576

Shivpuja, A. (2025). Transparency and trust building in data and AI: A framework for organizational success. *World Journal of Advanced Engineering Technology and Sciences, 15*(2), 2105–2113. https://doi.org/10.30574/wjaets.2025.15.2.0720

Stepper, J. (2015). Working out loud: For a better career and life. Ikigai Press.

Strong, D. M., Lee, Y. W., & Wang, R. Y. (1997). Data quality in context. Communications of the ACM, 40(5), 103–110. https://dl.acm.org/doi/10.1145/253769.253804

Thaler, R. H., & Sunstein, C. R. (2008). Nudge: Improving decisions about health, wealth, and happiness. Yale University Press. https://doi.org/10.1016/j.soscij.2008.09.003

Trist, E., & Emery, F. (1951). Organizational choice: Capabilities of groups at the coal face under changing technologies. Tavistock.

Trist, E. L., & Emery, F. E. (1951). Some social and psychological consequences of the longwall method of coal-getting. Human Relations, 4(1), 3–38. https://journals.sagepub.com/doi/10.1177/001872675100400101

Tsamados, A., et al. (2024). Human–AI governance and teaming. AI Ethics Journal, 15(2). https://link.springer.com/article/10.1007/s43681-024-00489-4

Tushman, M. L., & Scanlan, T. J. (1981). Boundary spanning individuals: Their role in information transfer and their antecedents. Academy of Management Journal, 24(2), 289–305. https://www.hbs.edu/faculty/Pages/item.aspx?num=2819

Valente, T. W. (2010). Social networks and health: Models, methods, and applications. Oxford University Press.

Valente, T. W. (2012). Network interventions. Science, 337(6090), 49–53. https://pubmed.ncbi.nlm.nih.gov/22767921/

Valente, T. W. (2025). Network interventions: Engaging social networks to accelerate diffusion, improve performance, and speed organizational change. Productivity Press.

W3C. (2009). SKOS Simple Knowledge Organization System reference. World Wide Web Consortium. https://www.w3.org/TR/skos-reference/

W3C. (2012). OWL 2 Web Ontology Language document overview (Second Edition). World Wide Web Consortium. https://www.w3.org/TR/owl2-overview/

W3C. (2014). RDF 1.1 concepts and abstract syntax. World Wide Web Consortium. https://www.w3.org/TR/rdf11-concepts/

Waller, M. J., Zellmer-Bruhn, M., & Giambatista, R. C. (2002). Watching the clock: Group pacing behavior under dynamic deadlines. Academy of Management Journal, 45(5), 1046–1057. https://www.jstor.org/stable/3069329

Wang, N. (2015). Applications of social network analysis in institutional research (AIR Professional File, Article 136). Association for Institutional Research. https://www.airweb.org/docs/default-source/documents-for-pages/reports-and-publications/professional-file/apf-136-2015-spring_applications-of-social-network-analysis-in-institutional-research.pdf

Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. Journal of Management Information Systems, 12(4), 5–33. https://www.tandfonline.com/doi/abs/10.1080/07421222.1996.11518099

Watts, D. J. (2011). Everything is obvious: Once you know the answer. Crown Business.

Watts, D. J., & Dodds, P. S. (2007). Influentials, networks, and public opinion formation. Journal of Consumer Research, 34(4), 441–458. https://academic.oup.com/jcr/article-abstract/34/4/441/1820236?redirectedFrom=fulltext

Weill, P., & Ross, J. W. (2004). IT governance: How top performers manage IT decision rights for superior results. Harvard Business School Press.

Wenger, E. (1998). Communities of practice: Learning, meaning, and identity. Cambridge University Press.

Wharton School of the University of Pennsylvania. (2019). HR management and analytics [Executive education program]. Wharton Executive Education.

Wilkinson, L. (2023, August 30). Walmart rolls out generative AI-powered assistant to 50K employees. CIO Dive. https://www.ciodive.com/news/Walmart-generative-AI-tool-My-Assistant/692385/

Kimura, Y, W. (2025). LEAD multi-organization behavioral telemetry dataset, 2020–2025 [Unpublished internal dataset]. LEAD.

Wu, Y. J., Antone, B., Srinivas, A., DeChurch, L. A., & Contractor, N. S. (2021). *Teamwork in the time of COVID-19: Creating, dissolving, and reactivating network ties in response to a crisis*. Journal of Applied Psychology, 106(10), 1483–1492. https://psycnet.apa.org/doiLanding?doi=10.1037%2Fapl0000969

Wu, Y. J., Antone, B., DeChurch, L. A., & Contractor, N. (2023). *Information sharing in a hybrid workplace: Understanding the role of ease-of-use perceptions of communication technologies in advice-seeking relationship maintenance*. Journal of Computer-Mediated Communication, 28(4), zmad025. https://academic.oup.com/jcmc/article/28/4/zmad025/7210242?login=false

Younis, S., Ahsan, A., & Chatteur, F. M. (2023). An employee retention model using organizational network analysis for voluntary turnover. Social Network Analysis and Mining, 13, 28. https://link.springer.com/article/10.1007/s13278-023-01031-w

Appendix A: LEAD Case Studies and Network-Based Organizational Insights

The appendix expands on illustrative evidence that supports the conceptual model, but is not intended as a full empirical analysis.

A1. Case Study 1

Findings from LEAD's dataset, spanning over 800 organizations between 2020 and 2022, reinforce a key theme from the literature: organizational learning and innovation depend less on formal hierarchy and more on the informal networks that sustain knowledge flow. Across deployments, employees who voluntarily initiated peer-connection or onboarding programs—often product managers, engineers, or team leads—served as de facto "trust brokers," meaning informal brokers who connect otherwise disconnected groups and enable knowledge sharing (Burt, 1992; Cross & Parker, 2004; Valente, 2010). Their behavior is consistent with prior studies on social capital and network diffusion, which find that knowledge sharing and innovation scale only when informal influence is supported, visible, and reciprocated (Burt, 1992; Cross & Parker, 2004; Valente, 2010).

When these employees lacked reinforcement such as leadership sponsorship, aligned budgets, or integration with formal systems, connection programs decayed. Roughly 80% of early adopters who initiated them showed one or more exit or attrition indicators within 90 days. This was inferred using two observable indicators:

- Approximately 80% of corporate email addresses began returning bounces, suggesting deactivation, redirection, or administrative changes.

- 48.99% were verified on LinkedIn as having transitioned to new jobs within that period. LinkedIn confirmation is incomplete due to self-reporting and update lag.

Triangulating these indicators provides a useful proxy for early-stage attrition, suggesting that when organizations fail to retain informal influencers, they not only lose individual talent but also risk dismantling trust networks essential for knowledge continuity. Retention is therefore not merely an HR metric but a structural determinant of AI and KM readiness, since knowledge, literacy, and behavioral consistency are less likely to compound when the stewards of informal systems depart.

Although this study does not compute firm-level financial loss for privacy reasons, the resolution of the dataset could support cost estimation in principle. Network research suggests that employees who occupy bridging or high-centrality roles generate disproportionate value by maintaining coordination pathways and enabling tacit knowledge flow, and their exit is associated with reduced information diffusion and team performance (Burt, 1992; Cross & Parker, 2004; Borgatti et al., 2018). Human-capital scholarship provides standardized turnover-cost components such as recruitment, onboarding, productivity loss, and erosion of tacit expertise, which allow for conceptual bounding without confidential data. Because collaboration metadata reveals dependency patterns, cross-functional exposure, and knowledge-flow centrality, and because labor-market benchmarks offer role-specific and region-specific salary proxies, triangulated indicators such as email deactivation signals and LinkedIn transitions can provide an analytical foundation for estimating the potential economic impact of losing these informal brokers. The purpose of this study is not to calculate these costs but to clarify that the erosion of informal networks represents a potentially quantifiable organizational vulnerability rather than a purely qualitative one.

A2. Case Study 2 Network Connectivity and Organizational Outcomes

Two organizations within LEAD's dataset illustrate how network structure and informal connectivity shape learning, engagement, and business-relevant outcomes. Both adopted LEAD.bot between 2021 and 2024, with observational data drawn from 2021 to 2023. Their trajectories diverged based on baseline culture, network maturity, and learning conditions.

Company A: Strengthening Fragmented Networks

Company A, a global enterprise with approximately 1,500 employees, illustrates how weak social connectivity can undermine knowledge reliability and organizational learning. When the company's HR team first engaged LEAD in 2021, they reported unusually high new-hire turnover and declining engagement scores—issues that predated the deployment and were attributed to fragmented collaboration and limited cross-team visibility rather than compensation or performance factors. LEAD was introduced as one of several interventions to strengthen informal networks and improve early-stage integration.

At the time of rollout, fewer than 1 percent of employees functioned as connectors, well below the 5 to 35 percent benchmark often used as a practical range for resilient, innovation-supporting networks (Wharton School of the University of Pennsylvania, 2019). . This relatively low level of informal connectivity coincided with limited peer learning and slow knowledge transfer between teams, which leaders described as weakening the organization's capacity to retain and onboard talent effectively. Several onboarding touchpoints, including structured introductions, shadow calls, and short in-team training rotations, offered the company practical ways to integrate new hires while simultaneously providing behavioral signals for LEAD. Patterns in how quickly newcomers were pulled into real work, how consistently they engaged with these touchpoints, and where participation stalled served as indirect indicators of onboarding effectiveness. In Company A, these signals aligned with the organization's high early attrition rate, revealing how weak informal pathways hindered new-hire integration.

Behavioral Nudges and Cultural Campaign

To strengthen informal networks, the HR team worked with LEAD to embed small, visible behavioral nudges into daily routines. These included:

- a company-wide “connection campaign,” with printed mugs, T-shirts, and a clear slogan explaining *why* the company was adopting LEAD and *how* employees were expected to use it
- recognition for employees who consistently met colleagues through LEAD.bot
- “lunch-and-learn” prompts and storytelling posts celebrating cross-team introductions

These nudges shifted interaction norms by making connectivity visible, socially endorsed, and behaviorally reinforced, consistent with behavioral-science mechanisms associated with nudging and norm activation (Thaler & Sunstein, 2008; Hallsworth & Kirkman, 2020).

Improvements in Knowledge Flow and Response Time

As these interventions took hold, connector density increased significantly by 2023. Patterns of knowledge transfer also improved. Average response time for internal requests decreased from approximately two days to under ten hours. Although high-performing teams typically aim for response times closer to two hours, this shift represented a substantial acceleration and materially improved cross-team coordination. Faster responses increased the likelihood that early-stage projects succeeded and reduced delays during onboarding, where timely support is critical.

Psychological Safety as a Moderating Factor

However, Company A also displayed low psychological safety, which constrained the durability of network improvements. A notable example occurred when a daylight-savings-related timing shift temporarily delayed LEAD.bot introductions by one hour. Although the technical issue was minor and quickly resolved, several IT and HR staff urgently escalated it, with one individual expressing fear of employment consequences if introductions were not sent precisely on time.

The reaction was not typical among LEAD’s customer base and reflected an environment where minor irregularities carried disproportionate perceived risk. Research in organizational behavior indicates that low psychological safety suppresses experimentation, discourages help seeking, and heightens the social cost of errors (Edmondson, 1999, 2019). In Company A, these dynamics amplified existing fragmentation and made behavior change more fragile.

Overall Outcome

Despite these constraints, strengthening the informal network yielded measurable benefits. Engagement and collaboration scores improved across departments, and leaders interviewed for this study noted a visible increase in employee morale, describing greater enthusiasm for cross-team projects and a stronger sense of shared momentum. Several employees who regularly participated in LEAD-enabled interactions later received internal innovation awards,

reflecting greater access to informal expertise and broader relational visibility. New-hire retention also improved during this period, which leaders attributed to faster integration, more reliable informal support, and clearer pathways for building early relational capital. Gains in project velocity and strategic execution were similarly linked to enhanced knowledge flow, better access to informal expertise, and more consistent cross-team responsiveness. This case demonstrates how reinforcing informal networks, supported by targeted behavioral nudges, can rebuild knowledge pathways in structurally fragmented environments, while also showing how limited psychological safety can moderate the stability and scalability of such improvements.

Company B: Scaling from High Connectivity

Company B, a consulting firm of approximately 3,000 employees, illustrates how strong baseline network connectivity and psychological safety accelerate the adoption of people-side KM practices. As a project-based organization, teams are routinely assembled and dissolved across engagements, requiring employees to collaborate with new colleagues, navigate expertise quickly, and form temporary but high-functioning networks inside client environments. This operating model creates a workforce already accustomed to rapid relational formation, high interdependence, and frequent cross-boundary knowledge exchanges.

When LEAD was first deployed in 2022 in a limited proof-of-concept within the marketing department, more than 40 percent of employees already exhibited connector behavior. Even those outside the pilot group demonstrated relational patterns aligned with company-wide cultural expectations: openness, initiative in cross-team interactions, and comfort navigating unfamiliar collaborators. As a result, employees intuitively understood the purpose of LEAD with minimal additional socialization or promotion, suggesting that the program benefitted from strong employee “preconditions”—a cultural environment in which network engagement was already normative and valued.

Within this context, internal response times for cross-team knowledge requests averaged ~2 hours—faster than what LEAD typically observes in comparably sized deployments. Coordination across functions showed high reciprocity and low reliance on formal escalation, based on consistent participation and timely follow-through in introduction workflows. Public employer ratings (e.g., Glassdoor) showed no immediate shift during the pilot period; however, after broader rollout, the organization’s external sentiment improved. Because public ratings are shaped by multiple factors and are not used here to attribute causal impact, this observation is presented as contextual rather than evaluative. The pattern is consistent with prior work

suggesting that reciprocity and recognition in informal networks can strengthen perceived fairness and belonging.

Psychological Safety as an Enabling Condition

Company B also demonstrated high psychological safety, which appeared to moderate how employees interpreted minor anomalies. A temporary timing discrepancy occurred during a regional daylight-savings shift—the same type of delay that triggered urgent escalation and fear of punitive consequences at Company A. In contrast, Company B's team neither escalated the issue nor viewed it as consequential. When informed, they simply acknowledged they had noticed the shift and expressed no concern.

This reaction aligned with broader behavioral patterns in the firm: high connector density, open cross-team collaboration, comfort with ambiguity, and low fear of minor process variation. Research shows that environments with higher psychological safety adopt new practices more smoothly and interpret anomalies as learning signals rather than threats (Edmondson, 1999; Schein, 2010). Company B's experience supports this pattern and indicates that psychological safety may function as a multiplier for people-side KM interventions, increasing the stability of behavioral change and accelerating benefits.

Overall Outcome

By combining a consulting-oriented operating model, strong baseline relational connectivity, and high psychological safety, Company B demonstrated how organizations with favorable preconditions can scale people-side KM tools with minimal friction. Employees engaged consistently, network pathways strengthened further, and the organization achieved visible cultural gains without altering formal incentives. This case provides a contrast to Company A by showing how preexisting social infrastructure shapes both the trajectory and speed of people-side KM adoption.

Comparative Insight: Network Maturity, Psychological Safety, and Voluntary Collaboration Capacity

Across the two cases, a consistent pattern emerges. Organizations differ not only in their structural network maturity but also in their voluntary collaboration capacity, meaning the degree to which employees engage across boundaries without formal obligation. This capacity, long emphasized in KM and organizational learning research, shapes how effectively new practices, tools, and ideas diffuse through the organization.

Company A illustrates how fragile informal networks can become in environments with low psychological safety. Even modest system anomalies, such as a daylight-savings-related

introduction delay, triggered urgent escalation and fear of reprisal. These reactions reflected weak cross-team ties, low connector density, and limited willingness to engage in discretionary collaboration. Although network strength improved through deliberate nudges such as lunch-and-learn sessions, public storytelling, and visible rewards, the initial learning environment constrained both response speed and early experimentation.

In contrast, Company B demonstrated strong baseline connectivity, high psychological safety, and a deeply embedded norm of cross-functional teaming. The same timing anomaly was treated as inconsequential, and employees continued to participate actively without additional prompting or reinforcement. This experience aligns with the firm's consulting-driven work model, where employees routinely assemble new teams, collaborate across functions, and integrate into client environments. These cultural preconditions supported rapid voluntary collaboration and reinforced strong knowledge-sharing habits.

A notable difference also appeared during the transition from pilot to full deployment. Both organizations conducted six-month pilots prior to their annual subscriptions, yet activation speed diverged sharply. Company A required nearly three months of visible internal promotion before employees began forming cross-team connections at scale. Company B reached near-universal participation within the first week, despite minimal promotion. Activation curves therefore function as behavioral signals of relational readiness: in network-mature, psychologically safe environments, employees treat connection programs as extensions of existing norms, whereas in fragmented or risk-averse settings, discretionary collaboration emerges only after prolonged socialization.

A separate insight from Company A highlights a different mechanism operating at the individual level. One employee who joined during the initial pilot period explained that early relational access gained through introduction and buddy-style interactions gave her visibility into tacit practices and organizational dependencies. These early ties helped her identify meaningful challenges in the organization and ultimately contributed to her advancement into management. She also observed that colleagues who actively participated in the knowledge-sharing program experienced similar benefits. This pattern aligns with proactive socialization theory, which shows that newcomers who actively build relational capital gain faster integration and upward mobility (Ashford and Black, 1996; Cooper-Thomas and Anderson, 2006), and with research showing that strong early social ties predict retention and performance (Cable, Gino, and Staats, 2013).

Role of Structured Socialization Mechanisms in Revealing Collaborative Capacity

A key observation across deployments is that structured socialization mechanisms, such as employee-introduction programs, and buddy programs, act as activation points that generate measurable evidence of an organization's collaborative capacity. Prior research shows that such mechanisms scaffold learning by creating repeated opportunities for interpersonal exchange, which can reveal whether employees choose to collaborate when no formal requirements exist (Morrison, 2002; Bauer and Erdogan, 2011).

Across deployments, these mechanisms contributed several forms of insight relevant to people-side KM and AI-readiness models:

• **Behavioral activation signals**
Standardized pairing or introduction moments allow organizations to observe who responds, who reciprocates, and which units opt into cross-boundary collaboration. These signals illuminate the underlying strength of informal knowledge pathways, consistent with network-based models of social learning (Valente, 2010).

• **Relational choice data**
Voluntary participation patterns provide insight into which relationships are likely to support meaningful knowledge exchange. Many organizations conduct internal surveys before adopting programs such as buddy systems or employee introductions, and these surveys often reveal a latent desire for cross-functional connection, mentoring, or informal learning. During early pilots, participation in LEAD-supported introductions was typically voluntary, and only a very small minority of employees opted out.

Within this voluntarily participating group, engagement varied. Relational choice signals distinguish employees who reliably act on their stated desire to connect from those who participate only intermittently or whose time-use patterns and ONA roles suggest interest but constrained capacity. These distinctions matter because they reveal not only structural opportunities for collaboration but the behavioral conditions that enable tacit knowledge to flow..

Voluntary participation rates, response delays, and repeated engagement patterns therefore reveal which ties function as reliable conduits for tacit knowledge. This complements structural ONA by showing not only who could collaborate but who consistently does, aligning with research on relational coordination and cooperative work, which emphasizes that consistent, reciprocated action is a stronger predictor of knowledge transfer than structural proximity alone (Gittell, 2002).

• **Contextual indicators of psychological safety and teaming norms**
Patterns of engagement reflect whether employees interpret new interactions as opportunities or risks. This aligns with psychological safety research showing that small behavioral cues indicate whether individuals feel free to learn, ask questions, or take interpersonal risks (Edmondson, 1999).

These insights are useful not because they belong to any specific system, but because they represent a class of data that traditional enterprise platforms rarely capture. Structured socialization moments function as organization-wide micro-experiments, revealing latent norms, trust structures, and collaboration pathways that are essential inputs for people-centered AI models.

Across the broader dataset, organizations with greater voluntary collaboration capacity, stronger cross-functional ties, and higher psychological safety realized faster improvements in responsiveness, knowledge flow, and innovation outcomes. Organizations with fragmented or risk-averse cultures changed more slowly, even when structural interventions were introduced. This reinforces a central conclusion from KM and organizational learning research: informal networks form the operational layer of knowledge flow, and their strength predicts retention, learning continuity, and adaptability (Davenport and Prusak, 1998; Cross and Parker, 2004; Edmondson, 1999).

Taken together, the comparative evidence shows that network maturity, psychological safety, and voluntary collaboration capacity jointly shape an organization's ability to absorb new knowledge, adopt emerging technologies, and maintain learning continuity during change. Organizations in which these dimensions are strong tend to move faster, learn more consistently, and exhibit measurable performance gains. This provides empirical support for people-side KM as foundational to sustainable AI adoption.

A3. Quantifying Knowledge Flow: From Behavioral Signals to Measurable Indicators

Across more than 1,500 organizations analyzed between 2020 and 2025, LEAD's aggregated dataset shows a consistent pattern: organizational performance, retention, and AI readiness are shaped less by formal structure and more by observable, repeatable behavioral dynamics. The two case studies discussed earlier—spanning fragmented to highly connected networks—illustrate that collaboration patterns, trust relationships, and responsiveness cycles can be quantified and used as decision-relevant indicators of organizational health.

The approach used in this study focuses on the behavioral layer of knowledge: how people actually connect, exchange information, allocate time, and coordinate work. Rather than emphasizing repositories or documented assets, it draws on several categories of privacy-preserving enterprise signals. Importantly, the system does not capture all possible data; instead, it applies sampling, aggregation, and privacy-preserving techniques to reduce intrusiveness, limit employee concern, and ensure that insights reflect collective patterns rather than individual actions.

1. Passive Organizational Network Analysis (ONA)

Passive ONA approximates the structural opportunities through which knowledge may flow. Conceptual examples include communication frequency and distribution, meeting co-attendance, shared collaboration spaces, cross-functional exposure, and structural cohesion. These signals describe the topology of everyday collaboration and highlight where knowledge can, in principle, travel.

2. Active ONA and Relational Metrics

Active ONA captures the relational qualities that determine whether structural pathways are actually used. Examples include trust assessments, perceived expertise, willingness to provide or seek advice, and indicators of psychological or relational safety. These signals reveal the motivational context that shapes tacit knowledge exchange.

3. Temporal and Time-Use Rhythms

Time-based metadata surfaces the pace and stability of knowledge movement. Patterns in meeting load, working-hour alignment, request turnaround times, and responsiveness cycles illustrate temporal bottlenecks and reveal behavioral latency—information that cannot be inferred from static network structure alone.

4. Workflow and Coordination Patterns

Aggregated workflow indicators capture how work progresses across boundaries. While not tracking individual actions, they illuminate where interdependencies slow execution, where tasks cluster, and how smoothly coordinated work actually moves. These operational patterns provide an important complement to structural and relational insights.

5. Interaction Stability and Relational Reliability

To approximate the reliability of knowledge pathways, the model incorporates indicators of interaction stability—whether teams maintain recurring, dependable engagement patterns or exhibit fragmentation and inconsistency. Stable relational rhythms generally support tacit knowledge flow; irregular patterns often signal weak or at-risk knowledge infrastructure.

6. Conceptual Translation into Measurable Indicators

To maintain confidentiality, the appendix does not describe mathematical transformations, model weights, feature-engineering logic, internal benchmarks, or algorithmic thresholds. Instead, it outlines conceptually how heterogeneous behavioral signals can be synthesized into high-level composite indicators that approximate relational reliability, cross-functional connectivity, and the flow of expertise.

Indicator (Conceptual)	Purpose	Interpretation
Network Density	Measures the overall connectedness of the communication graph within a given time frame	Reveals whether collaboration is siloed, clustered, or adaptive to change
Cross-Functional Collaboration	Assesses the share of relationships bridging distinct teams or functions	Higher proportions indicate broader knowledge circulation and innovation potential
Trust Reciprocity	Captures the balance of mutual interactions (e.g., bidirectional messages, consistent responses)	Reflects relational reliability and psychological safety within teams
Knowledge Latency	Tracks the time it takes for information or practices to diffuse across the network	Lower latency indicates faster learning, shorter feedback loops, and higher adaptability

As noted earlier in the literature review (Part 4.3), constructs such as trust density illustrate how LEAD must translate between passive and active ONA when synthesizing behavioral and relational signals into measurable indicators. Trust density combines observable interaction opportunities (passive edges) with explicitly confirmed relational qualities (active ties), providing an example of how heterogeneous inputs can be made analytically comparable. The same principle underlies other indicators described in this appendix: because passive and active ONA do not map cleanly onto one another in existing literature, software systems require derived constructs that bridge structural opportunity with relational motivation. These conceptual translations enable behavioral signals to be expressed in machine-usable and governance-compatible form while preserving their relevance to knowledge flow, influence, and organizational learning.

Another central example is knowledge latency, which in this study differs from structural graph measures such as diameter or shortest-path distance. Whereas those metrics describe formal topology, knowledge latency reflects observed temporal behavior: the time it takes for information, assistance, or decisions to move through real coordination rhythms. Similarly, while formulas are not disclosed, the behavioral model incorporates conceptual elements analogous to centrality, including bridging influence, timely access to expertise, and positional advantage, because these constructs are well established in diffusion and organizational learning research.

Team-Level Regularities, Social Constraint, and the Transformative Role of Influencers

Across deployments, LEAD's data revealed a pattern consistent with organizational-behavior and network research: individual behavior is shaped and often constrained by team-level norms. Interview and metadata patterns showed notable inconsistencies in interaction pace. Individuals sometimes moved more quickly in external communication than in internal exchanges, and in other cases waited for internal counterparts whose response timing was uncertain. These delays can arise from many sources, including differing priorities, misjudged urgency, or reliance on others for the next step. Yet the practical outcome is similar. Individual action becomes paced by the broader coordination environment rather than by personal intent, creating predictable latency at the team level even when individual intent is high, regardless of whether the limiting factor is structural, cognitive, or situational.

Prior studies show that teams develop shared temporal rhythms, expectations about responsiveness, and unwritten rules for coordination that influence how quickly work progresses, frequently overriding individual preferences or motivation (Waller, Zellmer-Bruhn, & Giambatista, 2002; Hackman, 2002; Edmondson, 1999). LEAD's observations align with this evidence. Within most units, employees exhibited similar turnaround cycles and knowledge-latency patterns, even when their personal tendencies differed. A small number of cultural champions or bridging contributors sometimes worked faster than their environment, yet their behavior remained bound by prevailing local norms, reflecting findings that structural and cultural context shapes individual action (Cross & Parker, 2004; Burt, 1992).

At the same time, the case evidence shows that these constraints are not fixed. When organizations identify the right influencers and reinforce their behaviors through targeted nudges, team-level dynamics can shift. Company A illustrates this interaction between constraint and leverage. Although the organization initially displayed weak informal networks and low psychological safety, the combination of structured introductions, occasional shadow calls, short in-team training rotations, public storytelling, visible recognition, and repeated nudges created an activation surface that allowed new joiners to become credible carriers of new relational patterns. Several employees who joined during the pilot period reported that early relational access helped them understand tacit practices, collaborate across functions, and identify meaningful organizational challenges. Their active participation and growing visibility influenced peers and gradually altered local norms. Over time, these reinforced behaviors contributed to measurable improvements such as response times decreasing from two days to under ten hours, increased connector density, and improved engagement scores, culminating in external recognition such as the company's first appearance on a Best Places to Work list.

This interaction between structural constraint and individual influence aligns with research on proactive socialization and behavioral diffusion. New employees who proactively build relational capital often gain faster integration and upward mobility, and their behaviors can shift group norms when supported by visible reinforcement (Ashford & Black, 1996; Cooper-Thomas & Anderson, 2006; Cable, Gino, & Staats, 2013). LEAD.bot, by creating repeated opportunities for low-stakes interaction, made these micro-patterns observable. This visibility allowed the organization to identify cultural champions early and amplify their impact, turning individual behaviors into team-level change even within a low-psychological-safety environment.

Interpretation and Contextual Use

These indicators are not interpreted uniformly across organizations. Behaviors such as responsiveness, meeting load, or cross-team interaction are not inherently positive or negative; their meaning depends on team norms, functional roles, business objectives, workload patterns, and cultural context. A product-development group may require rapid iterations, whereas research or compliance teams may progress through slower, more deliberative cycles. Effective interpretation therefore depends on context rather than universal benchmarks.

Privacy, Aggregation, and Organizational Patterns

LEAD's privacy-preserving methods identify patterns rather than individual behaviors, emphasizing structural opportunities for knowledge flow and relational reliability. Because indicators are modular and replicable, organizations can adapt them using their own workflow or survey data without accessing sensitive information.

Across deployments, several consistent patterns emerge: organizations with stronger cross-functional connectivity and higher trust reciprocity tend to exhibit lower knowledge latency, faster resolution of interdependent tasks, and more sustainable engagement over time. These findings show that social and behavioral dynamics—often treated as qualitative or difficult to operationalize—can be systematically observed, approximated, and incorporated into analytic models that support knowledge management, organizational learning, and readiness for AI-enabled work. Quantifying these dynamics provides a foundation for understanding how human networks evolve during digital transformation and how organizations maintain learning continuity amid technological change.

Appendix B: Interview Questions

[Name] thank you for interviewing with me today! My research is on effective Agentic AI adoption, and how behavioral insights (e.g., network structure, incentives, priming, messengers, etc.) affect the different success measures.

1. From your perspective, what are the biggest barriers to getting employees to use AI in their own work — for example, adopting tools like Copilot in daily workflows or using (or considering building) an internal knowledge assistant for engineering, customer support, or sales teams?

2. What has surprised you most in your organization's (or clients') internal AI adoption journey?

3. [Influence, Nudges, and AI literacy, and how these come about]

    3.1 My research is about these 3 questions, and I'd love your perspective on any or all:

Understanding influence — When new tech or AI is introduced, how do you notice who actually influences employees' adoption? (For example, "How are employees influenced by people who are in their networks, but are not specifically responsible for training those employees? Are there positive or negative influence examples that come to mind?)

Behavioral nudges — Do you use things like incentives, default options (what researchers call "choice architectures"), norms, or practices like hands on seminars to encourage people to try and adopt AI tools?

Building AI readiness / AI literacy — How do you help people know enough about AI to use it confidently and trust it? Do you have any plan or process for AI literacy?

3.2 Do you believe training is enough for adoption, or does it require a broader shift in business intelligence and organizational culture?

Appendix C: Practitioner Interviewees and Public Commentary Sources

1. Primary interviewees (practitioner perspectives)

- Christina Libs — Program Manager & Knowledge-Management Entrepreneur, [Company Redacted].
- Dr. Edward J. Hoffman — Former Chief Knowledge Officer, NASA; Founder & CEO, Knowledge Strategies LLC; Senior Lecturer, Columbia University. Founding Director of NASA's APPEL program and co-author of *The Smart Mission*.
- Eve Porter-Zuckerman — Independent Consultant; Former Chief Knowledge Officer, Isaacson, Miller. Former CKO at a leading executive-recruiting firm and consultant in organizational change and professional-services KM.
- Mihai Pohontu — Founder & CEO, Amber Studios.
- Nabeel Ahmad — Co-Founder & CSO, ChangeForce; Adjunct Professor, Columbia University, former Head of Market Development & Industry Learning at KPMG.

- [Name Redacted] — Tech Lead, Twilio.
- [Name Redacted] — Y Combinator-Affiliated Founder and Y Combinator Partner, [Company Name Redacted].

2. Public Commentary Sources Reviewed

- Altman, S. (2024–2025). AI & society: Keynotes and interviews. In B. Marr, *15 quotes on the future of AI*. Time Magazine. https://time.com/partner-article/7279245/15-quotes-on-the-future-of-ai/
- Bailey, P. (2024). Microsoft Future of Work Report summary [Tweet]. X. https://x.com/DynamicWebPaige/status/1751356222761296358
- Bevan, H. (2024, July 19). Post on tacit knowledge and learning networks [Tweet]. X. https://x.com/HelenBevan/status/1814197182578266127
- Behavioral Insights Team. (2025). *AI adoption: BIT report*. https://www.bi.team/wp-content/uploads/2025/09/BIT-AI-ADOPT-2025.pdf
- Blackmon, K. (2025). Why most AI adoption fails: Framework & fixes. LinkedIn Pulse. https://www.linkedin.com/pulse/why-most-ai-adoption-fails-framework-fixes-kelly-blackmon-cpcm--igt7c/
- Bornet, P. (2024, May 21). Post referencing IBM CEO survey [Tweet]. X. https://x.com/pascal_bornet/status/1792966274651398382
- Challapally, A. (2025). As cited in NANDA Initiative, *MIT just proved why 95% of AI projects are failing and how design leaders can fix it*. Medium. https://medium.com/design-bootcamp/mit-just-proved-why-95-of-ai-projects-are-failing-and-how-design-leaders-can-fix-it-b199c80ddf43
- Cheung, R. (2024, July 12). Post on guided implementation [Tweet]. X. https://x.com/rowancheung/status/1811810033845567930
- Cross, R. (n.d.). What is organizational network analysis (ONA)? https://www.robcross.org/what-is-organizational-network-analysis/
- Daher, D. (2025). *Protect human expertise* [Research report]. HFS Research. https://www.hfsresearch.com/research/protect-human-expertise/
- Fink, A. (2025, October 4). Who has the best people analytics team? [Podcast episode]. *Directionally Correct*. https://wrkdefined.com/podcast/directionally-correct/episode/145-alexis-fink-who-has-the-best-people-analytics-team-and-what-does-the-future-hold-for-alexis
- GrennAN, C. (2025, October 2). Post on behavioral transformation [Tweet]. X. https://x.com/conorgrennan/status/1973750801643217270
- Grant, A. (2025a). From best practices to “better practices”: Building skill sets fit for the future of work. World Economic Forum. https://www.weforum.org/stories/2025/02/skills-future-jobs-adam-grant/

- Grant, A. (2025b). ChatGPT did NOT title this podcast (w/ Allie Miller & Ethan Mollick) [Audio podcast episode]. *Rethinking with Adam Grant*. TED. https://www.ted.com/podcasts/rethinking-with-adam-grant/chatgpt-did-not-title-this-podcast-w-allie-miller-ethan-mollick
- Green, D. (2025, August 1). People analytics review: August 2025 [Post]. LinkedIn. https://www.linkedin.com/pulse/best-hr-people-analytics-articles-august-2025-david-green--y4hfe/
- Huang, J. (2025). AI and the future of work [Interview quotes]. In B. Marr, *28 best quotes about artificial intelligence*. https://bernardmarr.com/28-best-quotes-about-artificial-intelligence/
- Khan, F. (2025, March 5). How to succeed with AI adoption: Prosci report [Post]. LinkedIn. https://www.linkedin.com/posts/faisalkhanlinked_prosci-catalyst-report-the-human-side-of-activity-7303201312102301697-9sYl
- Klinghoffer, D. (2023, June 5). How a Microsoft HR people analytics exec became a star [Post]. LinkedIn. https://www.linkedin.com/posts/dawn-klinghoffer_people-analytics-is-often-viewed-as-an-hr-activity-7264601975613202432-QqP7
- Lalli, F. (2025, May 8). ROT: Return on trust in the AI era [Article]. Medium. https://medium.com/@fabiolalli/rot-return-on-trust-in-the-ai-era-defining-and-implementing-2009679d7545
- Leffer, N. (2025, October 7). As cited in T. DeGrange, post on starting small to build habits [Tweet]. X. https://x.com/TyeDeGrange/status/1975624214431584523
- Levie, A. (2025). Post on dynamic context for AI agents [Tweet]. X. https://x.com/levie/status/1962748473138401646
- Mohindra, A. (2022). Lightcast Connect conference speaker bio. IQPC. https://eco-cdn.iqpc.com/eco/files/event_content/iela-market-study-092724PwmReNr3dfhCS0eLOOBYZxo91TlTYvJCOYc7vH1S.pdf
- Nadella, S. (2025). The hardest part of AI is changing how people work. *Business Insider*. https://www.businessinsider.com/satya-nadella-microsoft-ceo-ai-changing-how-people-work-jobs-2025-6
- Nagarkar, N. (2025, October 7). Post on awareness-first adoption [Tweet]. X. https://x.com/nandeepn/status/1975681283427868684
- Napper, C. (2025, September 4). Recalibrating workforce decisions via people analytics & GenAI [Podcast episode]. *Elevate Your AIQ*. https://wrkdefined.com/podcast/elevate-your-aiq/episode/ep-83-recalibrating-workforce-decisions-via-people-analytics-and-gen-ai-with-cole-napper
- Ng, A. (2017, January 10). *Baidu's Andrew Ng on the future of artificial intelligence*. *TIME*. https://time.com/4631730/andrew-ng-artificial-intelligence-2017/
- Ng, A. (2019). *AI for everyone: Notes* [Slide deck]. DeepLearning.AI. https://static1.squarespace.com/static/5530dddfe4b0679504639dc1/t/647f1b56ec72aa1

b4f67046d/1686051676168/AI%2Bfor%2BEveryone%2Bnotes%2Bby%2BAndrew%2BNg.pdf
- Plunkett, J. (2024, September 3). How do we drive AI adoption? Seven lessons. Medium. https://medium.com/@jamestplunkett/how-do-we-drive-ai-adoption-057fc152b409
- ServiceNow. (2024). Challenges of generative AI adoption. *ServiceNow Blog*. https://www.servicenow.com/blogs/2025/challenges-generative-ai-adoption
- Tapadinhas, J. (2025, May 20). When AI becomes your teammate: MIT finds 60% gains lost without proper teaming. [Tweet]. X. https://x.com/jtapadinhas/status/1924738920090157530
- Torres, L. (2024a, September 28). Introducing our 5-step predictive AI adoption model. Medium. https://leopoldotorresazcona.medium.com/introducing-our-5-step-predictive-ai-adoption-model-28ee83396830
- Torres, L. (2024b, November 4). Shaping the future of hybrid human–AI teams through network science. Medium. https://leopoldotorresazcona.medium.com/shaping-the-future-of-hybrid-human-ai-teams-through-network-science-9b65e8e0fd09
- Ulrich, D. (2021, September 11). How to build excellence in people analytics [Podcast episode]. *Digital HR Leaders*. https://www.myhrfuture.com/digital-hr-leaders-podcast/dave-ulrich-how-to-build-excellence-in-people-analytics
- Vicsek, A. (2025, September 12). The new scarcity at work: Human connection. LinkedIn Pulse. https://www.linkedin.com/pulse/new-scarcity-work-human-connection-andras-vicsek-4cjwf