Please quote as: Grauer, T., Li, M. M. & Heidel, L. P. (forthcoming). Building the Runway Before the Plane: How Organizational History Shapes the Preparation for Agentic AI in Financial Services. 21st International Conference on Wirtschaftsinformatik (WI), : AIS eLibrary.

Building the Runway Before the Plane: How Organizational History Shapes the Preparation for Agentic AI in Financial Services

Research Paper

Teresa Grauer[1], Mahei Manhei Li[2], and Luc Philipp Heidel[3]

[1] University of Kassel, Information Systems, Kassel, Germany
teresa.grauer@uni-kassel.de
[2] University of St.Gallen, Information Systems, St.Gallen, Switzerland
mahei.li@unisg.ch
[3] Wavestone Germany AG, Hamburg, Germany
luc-phillip.heidel@outlook.de

Abstract. Agentic AI promises value for financial services. Yet none of the 16 organizations in our study has deployed a productive multi-agent system. All anticipate this shift. We therefore frame our study as an analysis of how regulated organizations prepare for agentic AI. Drawing on affordance theory and Leonardi's concept of imbrication, we analyze 16 interviews with senior leaders in German insurance companies and banks using Gioia methodology. We identify two patterns. In fragmentary imbrication, organizations deploy agents as standalone solutions, accumulating systems that do not compose. In anticipatory imbrication, organizations build shared infrastructure before multi-agent use cases demand it. The decisive factor is not technology but organizational history: prior transformations through Lean, Data Science, or Robotic Process Automation determine whether organizations can prepare proactively or must build everything at once. We contribute a process model of organizational preparation for agentic AI and four practical recommendations for CIOs and transformation leaders.

Keywords: *Agentic AI, Organizational History, Imbrication, Affordance Theory*

1 Introduction

Artificial intelligence (AI) is transitioning from assistive co-pilots to autonomous agents. Breakthroughs in reasoning, memory, and tool use have enabled a new class of AI systems that independently set goals, plan multi-step actions, and execute them with minimal human intervention (Holldack et al., 2026; Acharya et al., 2025; Feuerriegel et al., 2024). We define agentic AI as goal-oriented systems composed of one or more agents that continuously observe their environment, learn to use tools, and decide when and how to act (Baird & Maruping, 2021; Qu et al., 2025; Shavit et al., 2023). Financial services, with its data-rich, language-heavy operations and complex regulatory processes, is positioned to benefit from this shift. Industry investments in AI are projected

21st International Conference on Wirtschaftsinformatik,
September 2026, Linz, Österreich

to reach $97 billion by 2027, and 70% of banking leaders already report some level of AI engagement (Basso et al., 2025; Crisanto et al., 2024). Yet most financial services organizations currently operate at the level of individual GenAI tools, internal chatbots, retrieval-augmented generation (RAG)-based knowledge assistants or on single agent Proof-of-Concepts (PoC) and face fundamental questions about how to prepare their architectures, governance structures, and organizational designs for the transition to multi-agent systems (Singla et al., 2025). What makes agentic AI distinct from prior technology is a material property that matters for how organizations adopt it. Within this environment, organizations face a recurring choice between two approaches. They can deploy a single agent and resolve its constraints locally in order to deliver value quickly, or they can invest in the infrastructure, enterprise architecture, governance frameworks, and organizational design that enable multiple agents to collaborate as a system (Henfridsson & Bygstad, 2013). Each isolated deployment solves a local problem but creates material dependencies (Leonardi, 2011). Multi-agent systems are not collections of individual agents (Allmendinger et al., 2026). Our study reveals two patterns. In what we call fragmentary imbrication, organizations deploy each agent as a standalone solution, its own governance, its own data access, its own architecture. In what we call anticipatory imbrication, organizations invest in shared infrastructure before multi-agent use cases demand it. To understand these patterns, we draw on two theoretical perspectives. Affordance theory (Markus & Silver, 2008) helps explain what single agents and multi-agent systems enable and constrain and why the affordances of the system differ in kind from those of individual agents. Leonardi's (2011, 2013) concept of imbrication helps explain the process: how social agency (governance, routines, communication) and material agency (architectures, agent capabilities) layer into durable configurations and why prior layers shape what future layers are possible. No organization in our sample yet operates a productive multi-agent system.

We ask two questions. First, conceptually, how are the anticipated affordances and restrictions of multi-agent configurations expected to differ from those of single agents, and what does this difference imply for how technology and organization become interleaved (RQ1)? Second, how do financial services organizations prepare for the anticipated transition from single-agent deployments to multi-agent systems? (RQ2)

Based on 16 semi-structured interviews with senior leaders in German insurance companies and banks, using Gioia (Gioia et al., 2013), we make three contributions. First, the decisive difference between organizations that can prepare proactively and those that cannot is not technology but organizational history. Organizations that built process discipline and governance routines through earlier transformations, lean, data science, RPA, can prepare for multi-agent systems proactively, whereas those without such foundations must build everything at once. Second, in regulated settings governance prerequisites must be in place before deployment can begin; organizations that build governance, architecture, and stakeholder relationships proactively move faster, while those that wait become stuck. Third, we extend affordance reasoning to the anticipated affordances and restrictions of multi-agent configurations, connecting it to the digital-infrastructure and modularity literatures. Although our cases are drawn from financial services, the underlying mechanism, prior process and automation transfor-

mations, parallels transitions in industry and production, such as the cumulative adoption of Lean, Manufacturing Execution Systems (MES), and RPA. We therefore read our findings as a cross-industry account of organizational preparation for agentic AI. Based on these findings, we offer practitioners guidance on preparing for the transition.

2 Theoretical Foundations

2.1 The Evolution of AI in Financial Services

The integration of AI into banking and insurance is a cumulative evolution spanning nearly four decades, with the financial industry ranking among the earliest adopters of AI technologies (Herrmann & Masawi, 2022). Yet despite successful pilots, most institutions stall between "proving" and "scaling" phases, treating AI as a technological upgrade rather than systemic transformation (Haefner et al., 2023). Now the transition goes one step further. Co-pilot systems can analyze, recommend, and explain. Agentic AI is designed to plan, reason, and act. This represents a continuation of the long-term evolution of information processing in finance (Aldasoro et al., 2025). Each generation of AI has expanded the scope of tasks machines can perform. Agentic AI extends this scope to autonomous process execution and addresses the core limitation of previous waves by moving from insight generation to outcome generation (MIT Technology Review Insights, 2025). This shift changes how firms organize work. In co-pilot models, AI augments individual workers. In agentic models, AI becomes an organizational actor. The role of the human shifts from operator to supervisor. Humans manage a workforce of digital agents that handle the execution of banking services (Aldasoro et al., 2025). Regulatory frameworks in two jurisdictions shape the deployment context for agentic AI. The EU AI Act (Regulation (EU) 2024/1689) establishes a risk-based classification system where AI systems performing profiling of natural persons are always high-risk (Article 6), triggering mandatory human oversight requirements (Article 14). The GDPR further constrains automated decision-making through Article 22. FINMA Guidance 08/2024 requires institutions to assess how AI affects their risk profile, maintain centrally managed AI inventories with risk classifications, and ensure comprehensive documentation for material applications.

2.2 Affordance Theory in IS

Affordance theory originates in ecological psychology (Gibson, 1979) and has become a foundational lens for studying the relationship between IT artifacts and organizational outcomes in IS research (Markus & Silver, 2008). An affordance is a possibility for goal-oriented action that arises relationally between a technology's material properties and the objectives of a specified user group (Markus & Silver, 2008). Affordances depend on who engages the technology, for what purpose, and under what organizational conditions (Majchrzak & Markus, 2012). The same system affords different possibilities to different actors in different contexts. Affordances do not materialize automatically. They require actualization conditions, organizational properties

that must be present for a perceived possibility to become a realized outcome (Volkoff & Strong, 2013; Strong et al., 2014). The process unfolds in stages: perception of an action possibility, actualization through concrete organizational engagement, and the resulting organizational effect (Strong et al., 2014). Importantly, affordances and constraints are symmetric. A specific technical feature may afford action for one organization while constraining another, depending on their respective regulatory environments, capabilities, and goals (Majchrzak & Markus, 2012). Leonardi (2011) adds that affordances and constraints layer over time: as actors perceive possibilities and encounter obstacles, they modify practices, which reshapes conditions for subsequent perception and action.

One characteristic of current affordance theory, however, is its treatment of technology as a single, bounded artifact. The theory provides robust mechanisms for analyzing how one artifact generates affordances in relation to one set of users.

Work on digital infrastructure (Henfridsson & Bygstad, 2013), layered modular architecture, digital innovation (Yoo et al., 2010), and modularity (Baldwin & Clark, 2000) has long examined how compositional and emergent properties arise across interacting components. We build on this rather than claim a gap: our question is how affordance reasoning extends to the anticipated action possibilities that emerge when single agents are composed into multi-agent systems.

2.3 Imbrication: Technology and Organization Layer Over Time

General systems theory holds that a system can exhibit properties not derivable from its individual components, the whole exceeds the sum of its parts (Bertalanffy, 1968). This principle of emergence explains why multi-agent configurations differ from single agents. It does not explain how organizations manage this difference. For this, we draw on Leonardi's (2011, 2013) concept of imbrication. Leonardi distinguishes two types of agency. Material agency refers to what a technology does when humans engage it, not its static properties, but its active function (Leonardi, 2011; Kallinikos, 2012). Social agency refers to the goals, intentions, and routines through which humans organize work, governance processes, communication pathways, regulatory compliance practices, organizational structures. These two types of agency are analytically distinct but become interleaved through practice. Imbrication describes this interleaving as a gradual, iterative process. Organizations deploy technology (material agency). They adapt routines and governance to accommodate it (social agency). These adaptations create new conditions. New technology is deployed on those conditions. Each cycle layers social and material agency into a more durable configuration, like overlapping tiles forming a roof. Three properties make imbrication particularly useful as an analytical lens. First, materiality pre-exists action (Leonardi, 2013). The material infrastructure an organization encounters was not built for the purpose at hand. It was assembled over time for different reasons and now defines the boundaries of what is possible before any new organizational decision is made. Second, imbrication is cumulative. Each layer facilitates or constrains the next. A governance approval secured for one technology creates a precedent that reduces effort for the next. This resonates with Henfridsson and

Bygstad's (2013) finding that digital infrastructure evolves through generative mechanisms where each extension enables further extensions. Third, imbrication can involve level shifts. Most imbrication cycles are bounded: one technology, one set of organizational adaptations. But some transitions require a qualitative jump, from component-level to system-level integration, where the infrastructure built through prior imbrication cycles is necessary but not sufficient.

3 Research Approach

Given the exploratory nature of our research question we adopted a qualitative research design based on semi-structured interviews. The interview data was organized using Gioia, aiming to provide a clear and systematic approach to building theory from the interviewees' experiences while ensuring transparency and academic rigor (Gioia et al., 2013).

We conducted 16 semi-structured interviews with senior IT and business leaders in German-speaking financial services firms between October 2025 and January 2026. The sample comprises eleven insurance companies and five banks. We acknowledge that this sample is skewed toward insurance. We therefore treat banking-specific observations as indicative rather than representative and flag them as such where they inform our claims. Leveraging an established consultancy partnership, we engaged financial services organizations that are preparing for agentic AI and secured direct access to senior decision-makers. Interview partners held roles including CIO, Head of AI, Head of Data & AI, Solution Architect, AI Lab Lead, and Operations Lead. All of them either held budget responsibility for AI initiatives or had directly overseen generative or agentic AI projects. This level of access is important: the transition to agentic AI is shaped by strategic and governance decisions that are visible primarily at the executive level. We sought heterogeneity along two dimensions: organizational size and AI maturity stage. This variation allowed us to compare organizations at different points on the transition path and to identify patterns that hold across different starting conditions. Table 1 provides an overview of all informants. Each interview lasted approximately 57 minutes on average, with individual interviews ranging from 51 to 66 minutes. All interviews were recorded and transcribed with prior informed consent. Interviews were conducted in German. We developed a semi-structured interview guide following Mayring's (2015) problem-centric approach, structured around three thematic blocks: (1) concrete experiences with agentic AI projects, including architectural decisions and implementation outcomes; (2) organizational conditions, governance requirements, and scaling challenges; (3) strategic assessments, including competitive dynamics, vendor evaluation, and future deployment plans. The guide was designed to let informants describe their experience in their own terms before introducing theoretical concepts, consistent with the inductive logic of Gioia. Data analysis followed Gioia in three stages (Gioia et al., 2013). In the first stage, we derived first-order codes directly from transcripts using in-vivo coding, staying close to the informants' own language. This produced 69 first-order concepts. In the second stage, we grouped codes into sec-

ond-order themes, researcher-imposed conceptual categories that abstract from informant-level descriptions toward theoretical patterns. At this stage, the theoretical vocabulary from affordance theory and imbrication entered the analysis. We characterize this second-order step as abductive and theoretically informed, we asked not only what informants described, but how their descriptions related to material preconditions, affordance perception, and the layering of social and material agency. In the third stage, second-order themes were clustered into three aggregate dimensions. Coding was carried out by two researchers, who independently coded an initial set of transcripts. Disagreements were resolved through discussion until consensus, with a third author consulted in the few unresolved cases. Coding proceeded iteratively over several rounds. Saturation was observed: no new first-order concepts emerged from the final three interviews, and the set of 69 concepts and seven themes had stabilized before the last cases were coded. All interviews were conducted and coded in German. Quotations reported in this paper were translated into English by the authors and checked against the source transcripts by a second author; all quotations should therefore be read as translations from German.

Table 1. Role and experience of interviewee

ID	Industry	Position	Interview time (min)	Org Size (number of employees in GER)	Maturity of Organization
1	Insurance	Group Head IT Lab	57	6,000	Platform building
2	Insurance	Processmanager	51	13,000	Experimentation
3	Banking	Customer-Marketplaces	58	7,500	Platform building
4	Banking	Group Head Innovation Lead	52	17,000	Agentic workflows (production)
5	Insurance	Solution Architect	53	33,800	Single-agent
6	Insurance	CIO	62	7,000	Ideation
7	Insurance	Led Data Scientist	65	5,000	Experimentation
8	Insurance	Head of Data & AI	63	4,000	Agentic workflows (production)
9	Banking	AI Design Engineer	53	900	Agentic workflows (production)
10	Banking	Associate Director Datamanagement	54	4,000	Single-agent (production)
11	Insurance	Team Lead AI Services	54	600	Agentic workflows (production)
12	Insurance	Head of AI	53	5,000	Experimentation
13	Banking	Head of AI	58	4,000	Agentic workflows
14	Insurance	Head of IT Ops	66	33,000	Experimentation

15	Insurance	Team Lead IT	61	9,500	Ideation
16	Insurance	Head of AI Lab	59	12,000	Single-agent

4 Results and Contribution

Our analysis resulted in 69 first-order concepts, abstracted into 7 second-order themes, clustering into 3 aggregate dimensions. We present findings by aggregate dimension, illustrating each with representative evidence and cross-case comparison. For readability, Figure 1 presents an extract of the Gioia data structure.

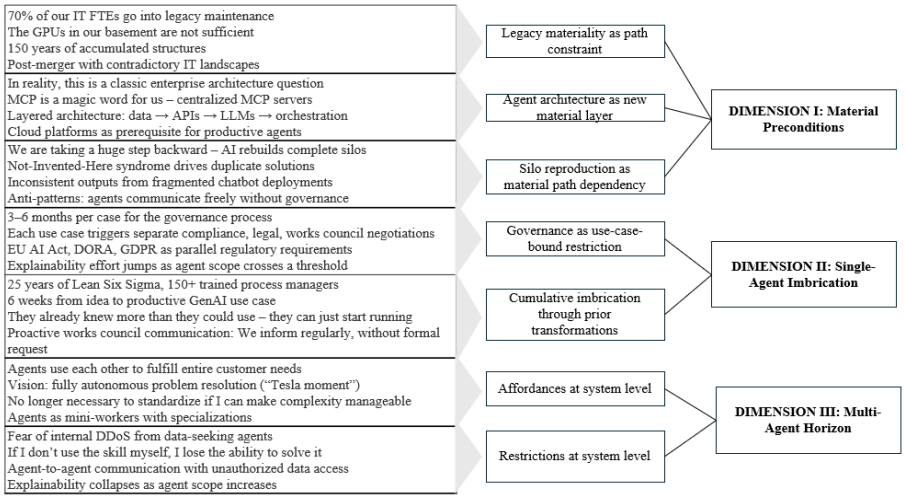

Figure 1.: Extract from Data structure

4.1 Dimension I: Material Preconditions

The first dimension captures the technical and organizational landscape that exists before any agent is deployed. It surfaces what Leonardi (2013) terms the materiality that pre-exists action, the infrastructure, systems, and architectural configurations that define the boundaries of what is possible before any organizational decision about agentic AI is made. Three themes constitute this dimension.

Theme 1: Legacy materiality as path constraint. Across all interviews, the pre-existing technical infrastructure emerged as a factor defining what organizations perceive as feasible. This pattern is particularly pronounced in banking. One informant estimated that approximately 70 percent of IT full-time equivalents are dedicated to maintaining legacy databases, systems built decades ago that no one modifies (I4). Another described "150 years of accumulated structures" (I10) constraining every modernization effort. Core transaction systems must process real-time payments against global sanctions lists within milliseconds. These systems will not be replaced. In the insurance sector, a recently merged organization operates with contradictory IT landscapes from two predecessor companies, creating friction that predates any AI initiative (I10). These

findings describe a materiality (Leonardi et al., 2012) that structurally conditions what comes next. A bank whose core runs on decades-old architectures faces different preconditions than an insurer that has migrated to cloud platforms.

Theme 2: Agent architecture as new material layer. All organizations that have moved beyond initial experimentation converge on a recognition: the transition to agentic AI is fundamentally an enterprise architecture question (I1). Informants describe a layered architecture: enterprise data at the base, retrieval-augmented generation and interfaces in the middle, interchangeable large language models above, and agent orchestration on top. Model Context Protocol (MCP) servers emerge as a central connectivity mechanism "a magic word for us" (I6) enabling agents to connect to core platforms without rebuilding them. One banking informant has built an automation platform processing 30 million documents annually and plans to replace its deterministic orchestration with agent-based routing (I4). This new architectural layer constitutes a qualitatively different materiality. Agents plan, access systems, and route decisions without continuous human initiation.

Theme 3: Silo reproduction as material path dependency. A counter-finding emerges from the data: the technology designed to integrate information silos reproduces them. Organizations use isolated chatbot solutions per department, each trained on local data, producing inconsistent outputs. "We are taking a huge step backward, because this AI makes it possible to rebuild complete silos. We buy tools that do really cool things on their own but only on their own" (I16). One informant observed a "Not-Invented-Here syndrome" (I5) driving duplicate solutions. Another warned against "anti-patterns where agents communicate freely without governance, creating quality and looping risks" (I2). This pattern illustrates how affordances and constraints layer over time (Leonardi, 2011). Existing material infrastructure decentralized procurement processes, absent central APIs affords departmental autonomy. Departments act on this affordance by purchasing standalone solutions. Each purchase adds a material layer, a new system with its own data model, its own governance approval, its own architectural assumptions that makes subsequent integration harder. The affordance of speed (deploy quickly per department) produces the constraint of fragmentation (agents cannot compose later) (I8, I12). It becomes visible only when the organization attempts to move from individual agents to multi-agent configurations.

4.2 Dimension II: Single-Agent Imbrication

The second dimension captures what happens when organizations deploy individual agents. It surfaces the iterative layering process, imbrication (Leonardi, 2011), through which social and material agencies become interleaved at the single-agent level.

Theme 4: Governance as use-case-bound restriction. Informants consistently describe governance as a dominant constraint on single-agent deployment. "We need three to six months per case to actually complete the governance process" (I14). Each agent deployment triggers separate negotiations with compliance, data protection, legal, and works council representatives. The EU AI Act requires proactive risk assessment. The Digital Operational Resilience Act (DORA) demands documentation for material applications. Works councils must be consulted on any technology affecting

working conditions (I2). These requirements apply independently per use case. The consequence is that governance functions as a social agency that cannot keep pace with the material agency of the technology. One informant attributed the delay not to incompetence but to the European regulatory fabric: "Everything I have done here in the last seven years, I could probably have done in China in seven months" (I4). A second informant offered a nuance: compliance itself is manageable through standard processes. The practical work of ensuring data quality, result reliability, and audit readiness for autonomous agents creates friction proportional to the agent's scope. Through the affordance lens, governance is not simply a barrier. It is a social structure that constrains certain affordances (speed of deployment, breadth of autonomy) while affording others (regulatory legitimacy, stakeholder trust, organizational learning about AI risks). The problem is that these governance affordances do not accumulate.

Theme 5: Cumulative imbrication through prior transformations. A contrasting pattern emerges in organizations with prior imbrication histories. One insurance organization deployed a productive GenAI use case within six weeks from idea to production (I12). The explanation lies not in technology but in organizational history. The organization introduced Lean Six Sigma 25 years earlier, training over 150 process managers. Data scientists were cross-trained with Lean experts. An internal knowledge platform was built for AI enablement. Communication pathways to the works council were established proactively: "We go regularly, tell them what we plan, without filing a formal request" (I2). The informant noted that their data scientists "already knew more than they could use. They can just start running." This case illustrates cumulative imbrication. Each prior transformation, Lean, data science capability building, RPA, created social and material layers on which subsequent layers could be built rapidly. The Lean infrastructure provided process discipline. The data science teams provided technical capability. The works council communication provided governance pathways. None of these layers was built for agentic AI. But together they constitute a dense fabric of interleaved social and material agency that absorbs new technology without requiring new organizational invention. The contrast with organizations that lack such histories is stark. Further organizations display partial anticipatory moves: two are building shared agent platforms and connectivity layers ahead of any concrete multi-agent use case (I1, I3), and one has institutionalized cross-functional AI enablement on top of an established data-science function (I8). Organizations sit on a continuum and can be fragmentary in their current deployments while building anticipatory infrastructure in parallel; several organizations in our sample (e.g., I4, I10) combine isolated departmental agents with early platform investments. We therefore read fragmentary and anticipatory imbrication as endpoints of a continuum along which organizations move over time, not as fixed organizational types.

4.3 Dimension III: Anticipated Multi-Agent Horizon

The third dimension captures the qualitative shift from single-agent to multi-agent configurations. General systems theory (Bertalanffy, 1968) holds that a system exhibits emergent properties when the whole possesses characteristics not derivable from its parts. It is irreducible: knowing everything about individual agents does not predict

what the multi-agent system will do. Our data show that practitioners anticipate this emergence, both its possibilities and its problems. No organization in our sample has deployed a productive multi-agent system. Yet all informants articulate visions of multi-agent futures and anticipate both the new action possibilities and the new restrictions. Two themes constitute this dimension.

Theme 6: Emergent affordances at the system level. Several informants articulate visions that exist only in multi-agent configurations. "An agent has a clear task and fulfills it. Agents use each other to fulfill entire customer needs" (I1). "By combining a variety of AI use cases, I can fully automate my processes. It is no longer necessary to standardize to the same degree if I can make complexity manageable" (I5). One informant envisions agents autonomously replacing hard-coded workflow engines in document processing pipelines handling 30 million documents annually (I14). These descriptions capture affordances that no individual agent produces. End-to-end process autonomy, complexity absorption without prior standardization, and self-organizing exception handling are properties of the configuration, not of any component. Informants recognize this. They describe multi-agent collaboration as different from deploying additional single agents. The perceived value lies "not in any single agent but in the compositional power of agent ecosystems" (I1). This perception corresponds to what general systems theory identifies as emergence: the whole possessing capabilities irreducible to those of its parts. Importantly, informants do not need the theoretical vocabulary to describe the phenomenon.

Theme 7: Emergent restrictions at the system level. Informants anticipate restrictions that do not exist at the single-agent level. One banking informant fears internal denial-of-service conditions from data-seeking agents overwhelming backend systems (I15). An insurance informant warns against agents communicating freely without governance, producing quality degradation and infinite loops. Identity and access management becomes critical, "which agent is allowed to access what, and how do we prevent an agent from obtaining data it should not have?" (I10). These restrictions are emergent in the sense that they arise from agent interactions, not from individual agent properties. A single agent accessing a database poses a manageable governance challenge. Multiple agents accessing multiple databases and triggering downstream actions based on each other's outputs poses a challenge that existing governance frameworks were not designed to address. The explainability that regulators require becomes progressively harder to provide as the number of interacting agents increases. One informant formalized this as a threshold effect: single-purpose tools pass through governance bodies with relative ease, but multi-function agents with broader scope require "categorically more explanation effort" (I12).

5 Discussion and Implications

Our findings reveal a transition that operates on two distinct levels. On the single-agent level, organizations deploy agents, negotiate governance, and adapt architecture incrementally. Each deployment is a bounded imbrication cycle. Governance frameworks

designed for individual agents no longer suffice. The challenge is not deploying agents. It is moving from one level to the other.

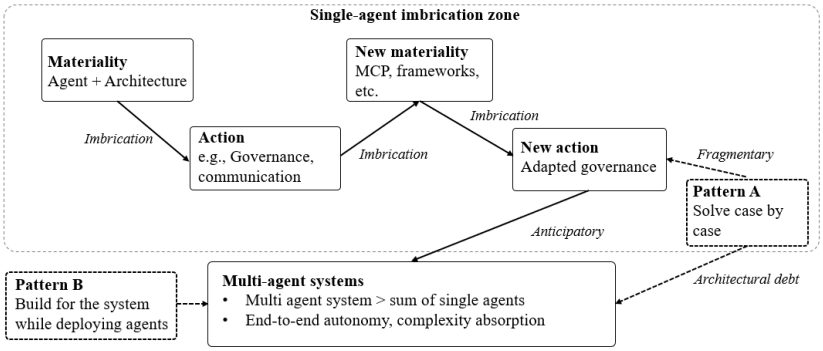

Figure 2.: Imbrication Transition Model

Figure 2 synthesizes our findings into a process model. The model reads from left (single-agent level) to right (anticipated multi-agent level). The horizontal axis represents time and accumulated imbrication; the vertical separation distinguishes the two preparation paths. Pattern A (fragmentary imbrication) shows isolated agents, each with its own governance and architecture, accumulating into non-composable silos and architectural debt. Pattern B (anticipatory imbrication) shows shared infrastructure, governance, and stakeholder communication built ahead of demand, so that the qualitative jump to multi-agent configurations becomes an incremental step. The dashed arrow marks the transition that organizational history makes either a leap (Pattern A) or a step (Pattern B). The decisive factor behind these patterns is not technology access. Both patterns appear in large, well-resourced organizations. The difference is organizational history. Each prior transformation created social and material infrastructure that reduces the cost of the next. These organizations can invest anticipatorily because they already have the organizational fabric to absorb the investment. Organizations without such history face a structural trap: they cannot justify building infrastructure without a proven multi-agent use case, but they cannot prove the use case without the infrastructure. This is the chicken-and-egg problem that keeps Pattern A organizations locked at the single-agent level. We offer three contributions to IS theory. First, and as our central contribution, organizational readiness for agentic AI is shaped less by current technology than by accumulated organizational history. Organizations carry forward the routines, communication pathways, and process discipline built during earlier transformations, and these prior layers determine how quickly new technology can be absorbed. We connect this path-dependence argument (Sydow et al., 2009) to research on absorptive capacity (Cohen & Levinthal, 1990) and organizational memory (Walsh & Ungson, 1991): readiness is a function of cumulative capability-building rather than of any single deployment. Second, IS research generally assumes that organizations adapt after technology is deployed (Volkoff & Strong, 2013). Our data suggest the opposite for regulated settings: governance prerequisites must be in place before deployment can begin, and organizations that build governance, architecture, and stakeholder relationships proactively move faster, while those that wait become stuck. Third, and most

modestly, we extend affordance reasoning from single technologies (Markus & Silver, 2008) to the anticipated affordances of multi-agent configurations. Emergence itself is not new: the digital-innovation and infrastructure literatures already theorize how layered-modular architectures and generative mechanisms produce properties no component possesses (Yoo et al., 2010; Henfridsson & Bygstad, 2013). Our addition is to read fragmentation through the same lens, agents built without shared design rules (Baldwin & Clark, 2000) cannot later be recombined, so the option value of modularity is lost before multi-agent use cases arise and to note that affordance perception here is anticipatory and governance-bound rather than realized.

Our findings yield four recommendations for CIOs and transformation leaders. First, most organizations deploy agents one by one. Each gets its own governance approval, its own data access, its own architecture. This produces working agents but agents that cannot work together. Organizations should ask: how do we build this agent so it can later collaborate with agents that do not yet exist? This means investing in shared infrastructure early, centralized APIs and common connectivity layers, even when only one agent uses them today. Second, the step from one agent to many is not about doing more of the same. When agents interact, new possibilities and new problems arise that single-agent governance does not cover. CIOs should plan for this explicitly. A practical test: what would our current governance framework cover if three agents worked together? What would it miss? Third, organizations do not start from zero. Lean, Agile, data science, and RPA programs leave behind trained people, established communication pathways, and process discipline. CIOs should inventory what earlier transformations left behind and activate these assets for the agentic AI transition, rather than building new enablement programs from scratch. Fourth, organizations that involve works councils and compliance only when a specific agent needs approval lose months each time. Those with ongoing dialogue move significantly faster. Organizations need governance partners who engage early and regularly, not only at the point of approval.

6 Limitations and Conclusion

Our study has several limitations. We interviewed AI leaders who drive the transition. Each organization is represented by a single senior informant, which exposes our account to single-informant bias; compliance officers, and works-council representatives, may assess the same challenges differently. No organization in our sample has deployed a productive multi-agent system. Our findings on multi-agent affordances and restrictions reflect what practitioners anticipate, not what they have experienced. Technical advances may narrow the governance gap between single-agent and multi-agent systems over time. Despite these limitations, the core finding holds across our sample: the transition to multi-agent AI is shaped less by technology than by what organizations built before. Those who invest in shared infrastructure early progress faster. Those who solve each agent in isolation get stuck. The transition may ultimately be decided not by who has the best agents, but by who builds the right layers.

References

Acharya, D. B., Kuppan, K., & Divya, B. (2025), Agentic AI: Autonomous Intelligence for Complex Goals: A Comprehensive Survey. IEEE Access, 13, 18912–18936.

Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2025), Intelligent financial system: How AI is transforming finance. *Journal of Financial Stability*, **81**, 101472. https://doi.org/10.1016/j.jfs.2025.101472.

Allmendinger, S., Bonenberger, L., Endres, K., Fetzer, D., Henner, G. & Kühl, N. (2026), Multi-agent AI. *Electron Markets* **36**, 18. https://doi.org/10.1007/s12525-025-00862-z.

Baird, A., & Maruping, L. M. (2021), The Next Generation of Research on IS Use: A Theoretical Framework of Delegation to and from Agentic IS Artifacts. *MIS Quarterly,* **45**(1), 315–341. https://doi.org/10.25300/MISQ/2021/15882.

Baldwin, C. Y., & Clark, K. B. (2000), Design Rules: The Power of Modularity. Cambridge, MA: MIT Press.

Basso, M., Belelieu, A., Hoang, J., Li, C., Propson, D., Parker, D., Picorelli, J., Premsundar, A., Smith, K., Vitolo, C., Krumpholz, K. W., & Caruso, J. (2025), Artificial Intelli-gence in Financial Services. World Economic Forum. https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf. Accessed 12.03.2026.

Bertalanffy, L. (1968), General System Theory: Foundations, Development. New York: George Braziller.

Cohen, W. M., & Levinthal, D. A. (1990), "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly* (35:1), pp. 128–152.

Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024), Regulating AI in the financial sector: Recent developments and main challenges. Financial Stability Institute. https://www.bis.org/fsi/publ/insights63.pdf. Accessed 09.03.2026.

Eidgenössische Finanzmarktaufsicht FINMA. (2024), *Aufsichtsmitteilung 08/2024: Governance und Risikomanagement beim Einsatz von künstlicher Intelligenz.* https://www.finma.ch/de/~/media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmitteilungen/20241218-finma-aufsichtsmitteilung-08-2024.pdf. Accessed 05.03.2026.

European Parliament and Council of the European Union. (2024), Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). EUR-Lex. https://eur-lex.europa.eu/eli/reg/2024/1689/oj. Accessed 05.03.2026.

Feuerriegel, S., Hartmann, J., Janiesch, C. & Zschech P. (2024), Generative AI. *Bus Inf Syst Eng* 66, 111–126. https://doi.org/10.1007/s12599-023-00834-7.

Gibson, J. J. (1979), The Ecological Approach to Visual Perception. Boston: Houghton Mifflin.

Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013), "Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology," *Organizational Research Methods* (16:1), pp. 15–31.

Haefner, N., Parida, V., Gassmann, O., & Wincent, J. (2023), Implementing and scaling artificial intelligence: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 197, 122878. https://doi.org/10.1016/j.techfore.2023.122878

Henfridsson, O. & Bygstad, B. (2013), The Generative Mechanisms of Digital Infrastructure Evolution. *MIS Quarterly*: *Management Information Systems*. 37. 907-931. 10.25300/MISQ/2013/37.3.11.

Herrmann, H., & Masawi, B. (2022), Three and a half decades of artificial intelligence in banking, financial services, and insurance: A systematic evolutionary review. *Strategic Change*, 31(6), 549–569. https://doi.org/10.1002/jsc.2525.

Holldack, F., Banh, L. & Strobel, G. (2026), Agentic information systems. *Electron Markets* 36, 5 https://doi.org/10.1007/s12525-025-00861-0.

Kallinikos, J. (2012), "MODELLING REALITY: CONTEXT, SYSTEM AND MEANING". ECIS 2012 Proceedings. 5m https://aisel.aisnet.org/ecis2012/5.

Leonardi, P. M. (2011), When flexible routines meet flexible technologies: Affordance,constraint, and the imbrication of human and material agencies . *Management Information Systems Quarterly* , vol. 35 (1), pp. 147–167.

Leonardi, P.M. (2013), When Does Technology Use Enable Network Change in Organizations? A Comparative Study of Feature Use and Shared Affordances. *MIS Quarterly*, 37, 749-775. https://doi.org/10.25300/MISQ/2013/37.3.04.

Leonardi, P. M., Neeley, T. B., & Gerber, E. M. (2012), How managers use multiple media: Discrepant events, power, and timing in redundant communication. *Organization Science,* 23(1), 98–117. https://doi.org/10.1287/orsc.1110.0638

Markus, M. L., & Silver, M. S. (2008), "A Foundation for the Study of IT Effects: A New Look at DeSanctis and Poole's Concepts of Structural Features and Spirit," *Journal of the Association for Information Systems* (9:10), pp. 609–632.

Majchrzak, A., & Markus, M.L. (2012), Technology Affordances and Constraints in Management Information Systems (MIS).

Mayring, P. (2015), Qualitative Inhaltsanalyse: Grundlagen und Techniken. 12th edn. Beltz Verlag, Weinheim Basel.

MIT Technology Review. (2025), Finding value from AI agents from day one. MIT Technology Review. https://www.technologyreview.com/2025/07/17/1119943/finding-value-from-ai-agents-from-day-one/. Accessed 08.03.2026.

Qu, C., Dai, S., Wei, X., Cai, H., Wang, S., Yin, D., Xu, J., & Wen, J. (2025), Tool Learn-ing with Large Language Models: A Survey. Frontiers of Computer Science, 19(8), 198343. https://doi.org/10.1007/s11704-024-40678-2.

Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O'Keefe, C., Campbell, R., Lee, T., Mishkin, P., Eloundou, T., Hickey, A., Slama, K., Ahmad, L., McMillan, P., Beutel, A., Passos, A., & Robinson D. (2023), Practices for Governing Agentic AI Systems. Research Paper, OpenAI, 1–23.

Singla, A., Sukharevsky, A., Hall, B., Yee, L., & Chui, M. (2025). The state of AI in 2025: Agents, innovation, and transformation. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai. Accessed 23.03.2026.

Strong, D. M., Volkoff, O., Johnson, S. A., Pelletier, L. R., Tulu, B., Bar-On, I., Trudel, J., & Garber, L. (2014), "A Theory of Organization-EHR Affordance Actualization," *Journal of the Association for Information Systems* (15:2), pp. 53–85.

Sydow, J., Schreyögg, G., & Koch, J. (2009), Organizational Path Dependence: Opening the Black Box. Academy of Management Review, 34(4), 689–709.

Volkoff, O., & Strong, D. M. (2013), Critical Realism and Affordances: Theorizing It-Associated Organizational Change Processes. *MIS Quarterly*, 37(3), 819–834. http://www.jstor.org/stable/43826002.

Walsh, J. P., & Ungson, G. R. (1991), Organizational Memory. Academy of Management Review, 16(1), 57–91.

Yoo, Y., Henfridsson, O., & Lyytinen, K. (2010), The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. Information Systems Research, 21(4), 724–735.