

Build the Medium: Why Organizational Intelligence Is Mechanism, Not Metaphor

May 2026

Witold Reichhart and Arnaud Gelas

Five-Paper Programme: Enterprise AI and Organizational Intelligence

This paper is part of a five-paper programme examining why enterprise AI fails in regulated environments, what architecture resolves it, and what emerges when that architecture operates at depth.

Paper A — Dynamics Blindness (Reichhart and Gelas) Diagnosis. Names the architectural failure mechanism: LLMs process tokens without tracing causal chains through organizational dependencies. Chain-of-thought, RAG, tool use, multi-agent systems, and scaling do not add the missing causal infrastructure.

Paper B — The Predictive Organization (Gelas and Reichhart) Architecture. Specifies the resolution: a tripartite structure — Map (state), Physics (dynamics), Player (agency) — coupling neural perception with symbolic reasoning, operating on claims-based knowledge with prevalence weighting.

Paper C — Build the Medium (Reichhart and Gelas) ← *this paper* Theory. Ten independent theoretical traditions converge on what organizational intelligence requires. Introduces the capability/fertility distinction and the autonomy-to-initiative transition as the real measure of AI maturity.

Paper D — Governed Intelligence Architecture (Gelas and Reichhart) Methodology. The practitioner companion: five-stage Governed Intelligence Lifecycle (Ingest, Consolidate, Curate, Expand, Apply) with governance requirements, epistemic immunity, and the path dependency argument for knowledge infrastructure investment.

Paper E — From Autonomy to Initiative (Reichhart and Gelas) Capstone. Three conditions for governed initiative. Graduated immersion systems as institutional analogue. Governance relocation mechanism. Six computational enrichments. Active inference as normative model. The domain graph as missing middle layer.

Causal spine: Enterprise AI fails because of dynamics blindness (A) → the resolution is architectural (B) → the architecture works because ten traditions converge on what living systems require (C) → the practitioner methodology is a governed intelligence lifecycle with epistemic immunity (D) → when the architecture runs at sufficient depth, it produces governed initiative — agents that perceive what matters through immersion, not instruction (E).

Abstract

Ten theoretical traditions — from cell biology to thermodynamics to social-systems theory — have independently arrived at the same conclusion about how collective intelligence works, without knowing it. When synthesized, they specify a coherent and mutually reinforcing set of conditions for organizational intelligence that no single tradition could derive alone, and that current enterprise AI strategy does not address.

Every organization already operates through a medium - the informal networks, mentoring relationships, institutional memory, and accumulated ways of working through which judgment propagates and new knowledge emerges. This medium is real but implicit: slow, fragile, confined to human interaction, and invisible to computational systems. We argue that the central task of enterprise AI strategy is to make this exist-

ing medium explicit, persistent, and available to computational agents - creating an infrastructure where tacit organizational patterns propagate without full extraction, where agents develop judgment through immersion rather than instruction, and where the organization itself becomes a cognitive system. We propose the Capability x Fertility matrix, showing that current AI maturity frameworks measure capability while systematically underinvesting in fertility - the organizational capacity to generate new value from existing knowledge. Path dependency makes the sequence of investment consequential: organizations that optimize for capability before building fertility face progressively higher costs to add the missing foundation.

Drawing on autopoiesis (Maturana and Varela), scale-free cognition (Levin), the free energy principle (Friston), distributed cognition (Hutchins), social systems theory (Luhmann), absorptive capacity (Cohen and Levinthal), knowledge creation theory (Nonaka and Takeuchi), the adjacent possible (Kauffman), dissipative structures (Prigogine), and hemispheric analysis (McGilchrist), we demonstrate that organizational intelligence is mechanism, not metaphor - the same phenomenon as individual intelligence operating at a different scale. We establish the fertility-first design principle, introduce the distinction between autonomy and initiative as the real measure of AI maturity, and identify the categorical limits of computation that make human participation a permanent design requirement.

This is Paper C of a five-paper programme. Paper A (Reichhart and Gelas, 2026a) diagnoses dynamics blindness as the specific failure mechanism in enterprise AI. Paper B (Gelas and Reichhart, 2026b) specifies the medium's architecture in engineering terms - the tripartite structure of state, dynamics, and agency layers operating on a claims-based knowledge representation with provenance, temporal decay, and competing interpretations. This paper provides the theoretical foundations that explain why that specific architecture is required and what its categorical limits are. Paper D (Gelas and Reichhart, 2026d) specifies the practitioner methodology — a governed intelligence lifecycle with epistemic immunity. Paper E (Reichhart and Gelas, 2026e) identifies the conditions under which the architecture produces initiative rather than mere autonomy.

Keywords: organizational intelligence, tacit knowledge, autopoiesis, distributed cognition, dissipative structures, collective intelligence, absorptive capacity, AI maturity, knowledge infrastructure, fertility, adjacent possible

1. Introduction

A cell biologist studying bioelectric signalling in frog embryos, a physicist modelling dissipative structures far from equilibrium, a cognitive scientist observing naval navigation teams, and a sociologist analysing how organizations maintain themselves through communication have each, independently, described the same phenomenon: collective intelligence that arises from the coordination substrate between components, not from the sophistication of any individual component. They do not cite each other. They work in different fields, use different methods, publish in different journals. Yet their findings converge on a single architectural requirement — and that convergence has direct consequences for how organizations should deploy AI.

An organization that is intelligent in this sense — that perceives patterns no individual could recognize, retains knowledge no single person carries, and generates responses beyond any one department's design capacity — operates differently from one that merely coordinates smart individuals. The difference is architectural. And the architecture has theoretical requirements that ten independent traditions specify with surprising precision.

Ten theoretical traditions converge on what those requirements are. Each addresses a different question about how collective intelligence works. Maturana and Varela's (1980) autopoiesis identifies what kind of system must exist: self-producing, maintaining identity through continuous operation. Levin's (2019, 2022) scale-free cognition explains why collective intelligence works at all: the same principles operate at every scale from cells to organizations, and investment in the coordination mechanism - the cognitive glue - matters more than investment in individual components. Friston's (2010) free energy principle provides mathematical formalization: systems maintain themselves by minimizing surprise through continuous model updating, and the boundary through which they sense and act - the Markov blanket (Kirchhoff et al., 2018) - determines their cognitive reach. Hutchins's (1995a, 1995b) distributed cognition demonstrates the empirical precedent: naval navigation teams exhibit system-level intelligence that no individual crew member possesses. Luhmann's (1995) social systems theory shows how social systems maintain themselves through communicative closure rather than physical boundaries. Cohen and Levinthal's (1990) absorptive capacity establishes the preconditions: an organization's ability to learn depends on its prior related knowledge. Nonaka and Takeuchi's (1995) SECI model maps how knowledge converts between tacit and explicit forms across individuals and groups. Kauffman's (2019) adjacent possible defines what can emerge from a fertile system - outcomes that cannot be specified in advance. Prigogine's (1980; Prigogine and Stengers, 1984) dissipative structures explain why the system requires continuous energy flow or it collapses. McGilchrist's (2009, 2021) hemispheric analysis establishes what computation cannot supply - and why human participation is not a transitional concession but a permanent design requirement.

These frameworks do not address the same question. They cannot be merged into a single grand theory. What they do is converge on a single architectural requirement: the informal medium through which every

organization already operates - its knowledge networks, judgment-propagation mechanisms, and institutional memory - must become an explicit, persistent infrastructure through which all ten mechanisms can operate together. We call this infrastructure the medium - what Levin (2019) terms the cognitive glue in biological systems - not because it is new, but because making it visible and formalizable is the precondition for organizational intelligence at scale. The medium is cognitive infrastructure, not data infrastructure. It is the substrate through which the organization coordinates, learns, and thinks.

The practical implications of this convergence are spelled out in companion papers. Paper A diagnoses the specific failure mechanism - dynamics blindness - that explains why current enterprise AI deployments underperform systematically (Reichhart and Gelas, 2026a). Paper B specifies the tripartite architecture required to build the medium: a state layer, a dynamics layer, and an agentic layer operating on a claims-based knowledge representation (Gelas and Reichhart, 2026b). This paper establishes the theoretical foundations. It explains why such an architecture is necessary, what it must satisfy, and where its categorical limits lie.

This is a theory paper. It does not present empirical results. It constructs a theoretical framework - grounded in established science across multiple disciplines - that explains why current AI adoption patterns produce the value gap the market data reveals, and what alternative infrastructure would be required to close it. The framework generates testable propositions and specific architectural requirements. Empirical validation is the necessary next step, and we outline a research agenda in the discussion. The contribution is the demonstration that ten independently established theoretical findings, when synthesized, specify a coherent and mutually reinforcing set of conditions for organizational intelligence - conditions that current enterprise AI strategy does not address.

The causal spine of the argument runs as follows:

Current AI maturity models measure capability but not fertility (Section 2) → this matters because tacit knowledge - the source of organizational judgment - cannot be extracted, only propagated through a medium (Section 3) → ten theoretical traditions specify what that medium must be: self-producing, scale-free in coordination, continuously predictive, distributed, communicatively closed, grounded in prior knowledge, capable of knowledge conversion, expansive of possibilities, sustained through flow, and permanently inclusive of human participation (Section 4) → the real measure of maturity is not autonomous execution but contextual initiative, which only develops through immersion in an active medium (Section 5) → the medium is a living system requiring continuous energy flow; without it, knowledge decays toward entropy (Section 6) → computation has categorical limits that make human participation a permanent design requirement, not a transitional concession (Section 7) → path dependency makes the sequence of investment consequential: capability-first architecture progressively forecloses the option to add fertility later (Section 8).

Each section advances one link in this chain and maps to the theoretical framework that supports it. Section 9 discusses contributions, limitations, and the research agenda.

2. The Maturity Model Gap and the Fertility Principle

Every major AI maturity framework measures capability - how much AI can do, how autonomously it operates, how deeply it integrates into organizational processes. Gartner, McKinsey, and MITRE (2024) each track multiple capability dimensions. These frameworks acknowledge that maturity is multi-dimensional - governance, culture, data quality, and organizational readiness all appear. Yet these conditions play supporting roles. The conditions that make capability productive get acknowledged in the framework then brushed aside in the investment. McKinsey identifies organizational readiness as the bottleneck; budgets flow to technical capability.

The evidence confirms this pattern. Only 7% of organizations have fully scaled AI enterprise-wide, and only 17% attribute more than 5% of EBIT to generative AI (McKinsey, 2025). 42% of companies scrapped most of their AI initiatives in 2025, up from 17% the prior year (S&P Global, 2025). Gartner projects that over 40% of agentic AI projects will be canceled by 2027 (Gartner, 2025). The failure data alone does not reveal whether the bottleneck is capability or environment. The diagnostic data does. McKinsey's 2025 survey found that workflow redesign - not model selection or technical sophistication - has the single biggest effect on whether organizations see bottom-line impact from AI. Gartner projects that through 2026, organizations will abandon 60% of AI projects specifically because they lack AI-ready data (Gartner, 2025b). Only 12% of organizations report data quality and accessibility sufficient for AI deployment (Precisely and Drexel University, 2025). The models work. The organizational environment does not support them.

Data readiness and workflow redesign are necessary but not sufficient. They solve the deployment problem - getting AI into operation. They do not solve the generativity problem - whether the organization learns from that operation, discovers patterns no one specified, or produces knowledge it did not previously possess. An organization with clean data and well-designed workflows can execute efficiently. It does not follow that it gets smarter through execution. The maturity frameworks stop at deployment. The harder question is what happens after.

This points to a dimension that current frameworks do not measure: fertility - the organizational capacity to generate new value from existing knowledge, not just execute existing capability more efficiently. The distinction has precedent. March (1991) identified the tension between exploitation (refining existing capabilities) and exploration (generating new possibilities) as the central adaptive challenge for organizations, and showed that systems biased toward exploitation outperform in the short run but lose adaptive capacity over time. Argyris and Schön (1978) drew a parallel distinction between single-loop learning (correcting errors within existing frames) and double-loop learning (questioning the frames themselves). Capability investment maps to exploitation and single-loop learning. Fertility investment maps to exploration and double-loop learning. The organizational learning literature has diagnosed this tension for

decades. What it has not provided is an architectural framework for resolving it in the context of AI infrastructure — which is what the medium is designed to do. The distinction between capability and fertility maps onto a matrix with four positions:

	Low Fertility	High Fertility
High Capability	Powerful Tools	Living Medium
Low Capability	Early Stage	Fertile Ground

Early Stage: organizations beginning AI adoption with neither mature capability nor fertile knowledge infrastructure. **Fertile Ground:** organizations that have built conditions for knowledge flow and pattern propagation but lack mature AI capability - they learn well but execute inconsistently. **Powerful Tools:** organizations with sophisticated AI systems operating in sterile environments - they execute efficiently but do not learn from operation. **Living Medium:** organizations where high capability operates within fertile infrastructure - they get smarter by operating.

The matrix is not freely navigable. Path dependency makes the sequence of investment consequential. The phenomenon is well-established: Arthur (1994) showed how increasing returns lock technologies into dominant positions regardless of efficiency, and David (1985) demonstrated the mechanism empirically with the QWERTY keyboard — a suboptimal standard that persists because switching costs compound faster than dissatisfaction. Sydow, Schreyögg, and Koch (2009) extended this analysis to organizational path dependence and identify the mechanism precisely: organizational path dependence moves through three phases - preformation (contingent early choices), formation (self-reinforcing dynamics), and lock-in (rigidity). Capability-first AI investment follows this trajectory. In the preformation phase, the choices look innocuous: the organization designs agent workflows around explicit goal specification, builds data pipelines optimized for retrieval speed, and structures feedback loops to measure task completion rates. Each decision is locally rational. But collectively they encode an assumption - that the purpose of the infrastructure is to execute instructions efficiently. In the formation phase, this assumption becomes self-reinforcing. Workflows designed for autonomous execution treat knowledge as input to be consumed, not as something generated through operation. The data architecture captures what agents need to act but not the contextual traces - decision provenance, relationship patterns, outcome histories - that fertility requires. Each optimization makes the next optimization along the same axis cheaper and any deviation more expensive. By lock-in, the organization has built an architecture where fertility has no place to attach. The route from Powerful Tools to Living Medium remains possible, but it now requires retrofitting an entire infrastructure that was designed around a different premise. Section 8 develops the cost curve this creates.

There is precedent in developmental biology. Certain developmental windows narrow over time. You cannot grow wings after the embryonic period ends - not because the genetic potential is absent, but because the architecture has committed to a different structure. Kauffman's adjacent possible formalizes

this: what is actual now determines what is next possible. An organization locked into capability-first architecture can still add dashboards, deploy new agents, and optimize existing workflows - the adjacent possible within the current paradigm remains large. What it cannot easily do is restructure the knowledge representation that sits beneath all of them. The adjacent possible outside the paradigm has narrowed. An organization that built fertility first faces no equivalent constraint.

The concept requires precision. Fertility is the organizational rate at which existing knowledge generates new knowledge, new capability, and new possibilities. It is not the amount of knowledge the organization possesses but the rate at which that knowledge flows, transforms, and expands what the organization can do. We identify six measurable dimensions:

Knowledge flow rate - how rapidly insights propagate across organizational boundaries. In low-fertility organizations, valuable patterns are discovered and lost because they never reach those who could act on them. Contextual richness - the degree to which knowledge includes the conditions under which it applies, the boundary cases, the history of prior attempts. Pattern diversity - whether the knowledge infrastructure maintains competing interpretations or converges prematurely on consensus. Every large organization contains tribal knowledge - “risk” means something different to Legal than to Trading, and those differences often reflect real contextual variation rather than errors to be resolved. The medium must hold competing definitions without forcing premature convergence. Temporal depth - how far back organizational memory extends and whether the infrastructure captures how knowledge has evolved. Integration maturity - how well knowledge from different domains combines into novel configurations. Governance completeness - whether formal governance covers the full knowledge lifecycle from creation through deprecation.

These dimensions connect to but are distinct from Cohen and Levinthal’s (1990) absorptive capacity - an organization’s ability to recognize, assimilate, and apply new external knowledge. The relationship requires precision, because without it fertility looks like a rebrand of an established construct.

Absorptive capacity predicts an organization’s ability to learn from external knowledge. It is intake-oriented: can the organization recognize, assimilate, and apply what comes from outside? The mechanism is prior related knowledge - you can only absorb what you have the existing base to recognize. Fertility predicts something different: the organization’s capacity to generate novel knowledge endogenously from its own operational traces. It is generativity-oriented: does the organization produce new patterns, new combinations, new possibilities through operation itself - not through importing external knowledge but through the recombination and propagation of what it already knows in ways that expand the adjacent possible?

The distinction is testable. An organization with high absorptive capacity and low fertility will rapidly adopt external innovations but generate few of its own. An organization with high fertility and low absorptive capacity will produce internal novelty but struggle to integrate external advances. The medium requires both, but the path dependency argument applies specifically to fertility: absorptive capacity can

be developed at any stage by hiring domain expertise and building learning infrastructure, but fertility requires architectural decisions about how knowledge is represented, connected, and allowed to recombine - decisions that become progressively harder to retrofit once capability-first infrastructure has hardened.

Cohen and Levinthal showed that absorptive capacity is cumulative: it depends on prior related knowledge. This is the precondition. Fertility is what happens when that prior knowledge base operates within infrastructure designed for generativity rather than retrieval. The six dimensions above are the specific architectural features that determine whether accumulated knowledge generates new knowledge or merely gets retrieved more efficiently.

The strategic implication is that the sequence of investment matters more than the quantity. Fertility should precede or accompany capability, because capability without fertility produces powerful tools while fertility with capability produces an organization that compounds advantage through operation. The theoretical basis for why this is so requires examining what organizational intelligence actually is - and that examination begins with the problem that has defeated knowledge management for three decades.

3. The Tacit Knowledge Problem

“We can know more than we can tell.” Michael Polanyi (1966) articulated this, and it remains the central unsolved problem of organizational knowledge. Every knowledge management system since has attempted to prove him wrong - extract, codify, document, store. The approach captured what could be articulated but lost what made experts expert. The tacit dimension - pattern recognition, contextual judgment, the intuitive sense for what matters - resisted extraction because it was never explicit in the first place. Knowledge management initiative failure rates of 50-70% are widely reported across the literature (see Dalkir, 2011 for a review). Davenport and Prusak (1998) documented how organizations systematically confuse information management with knowledge management — building repositories that capture documents while losing the contextual relationships that make knowledge actionable. Collins (2010) sharpened the analysis by distinguishing three forms of tacit knowledge — relational, somatic, and collective — and showing that only collective tacit knowledge (the kind embedded in social practices) resists individual articulation in principle, not merely in practice. The leading causes of KM failure are instructive: they describe systems that capture information but fail to maintain the conditions under which knowledge actually propagates.

Current AI approaches have not moved the bottleneck. Retrieval-augmented generation, expanded context windows, and explicit instructions are powerful at working with explicit knowledge. They retrieve documents, synthesize written information, follow detailed procedures. But the constraint remains: you can only give agents what you can articulate. The technology has changed. The limitation has not.

The medium dissolves this problem differently. Instead of extracting tacit knowledge, it creates conditions

where tacit patterns propagate through behavioral observation rather than verbal articulation. The mechanism works as follows. A person operates within the medium - makes decisions, handles exceptions, responds to situations. Tacit patterns leave observable traces: not articulations but structural imprints in the form of decision sequences, contextual associations, provenance chains, temporal patterns, and relationship structures. These traces are not themselves tacit knowledge. They are explicit representations of behavioral patterns that reflect tacit knowledge. The distinction matters. The medium does not capture the expert's intuition directly. It captures the downstream effects of that intuition in a form rich enough that agents and other humans can develop analogous pattern recognition through exposure.

The consulting pyramid illustrates the existing medium in action. It was a slow-motion medium for tacit knowledge transfer. Juniors developed judgment not through manuals but through years of exposure to decisions, consequences, and organizational patterns. The pyramid provided what instruction could not: the context that turns competence into judgment. It validated that exposure - not articulation - is the mechanism through which judgment develops. The pyramid had severe limitations: it took decades, worked only for humans, was available only to those with access, and was mediated entirely through interpersonal interaction. The medium compresses the temporal dimension and extends the mechanism to computational agents.

Nonaka and Takeuchi (1995) provided the framework for understanding how this works. Their SECI model identifies four modes of knowledge conversion: socialization (tacit to tacit, through shared experience), externalization (tacit to explicit, through dialogue and reflection), combination (explicit to explicit, through synthesis), and internalization (explicit to tacit, through learning by doing). They also introduced the concept of *ba* - a shared context where knowledge creation occurs. Without *ba*, the SECI spiral stalls. The medium is a technologically enabled *ba* - a shared context where all four conversions occur continuously, at scale, and with computational agents as participants alongside humans.

The critical difference from previous knowledge management is directionality. Traditional KM flows from human to system to storage. In the medium, knowledge flows in all directions simultaneously: agents observe behavioral traces left by experts (socialization), the system surfaces correlations the expert herself could not articulate (externalization), relationship-preserving knowledge structures synthesize patterns across domains (combination), and humans absorb what the system discovers, embedding it into their own judgment (internalization). This is the SECI spiral operating at organizational scale and at machine speed - and it is what makes the medium a knowledge-creating infrastructure rather than a knowledge-storing one.

Precision about limits is essential. The medium dramatically expands the boundary between what can and cannot be captured. It does not dissolve Polanyi's paradox. Paper B (Gelas and Reichhart, 2026b) develops the practical implication of this boundary by distinguishing three categories of invisible knowledge: tacit knowledge (irreducibly below conscious articulation), undocumented explicit knowledge (articulable but never written down - the exception-handling logic in SMEs' heads), and dark data (already

written but structurally invisible to computation). Each requires a different intervention: the medium for tacit patterns, structured elicitation for undocumented expertise, and extraction engineering for dark data. The three-category distinction matters because most organizations conflate all three and apply the wrong intervention to each.

Some knowledge remains irreducibly tacit - the embodied judgment, the contextual sensitivity, the feel for when standard procedure no longer applies. Section 7 establishes why this residual tacit dimension is categorically beyond current computation and why human participation is therefore a permanent design requirement, not a transitional concession. Human expertise becomes more valuable, not less - the medium handles the dimensions of knowledge work that can be computed, concentrating human judgment on the dimensions that cannot.

4. Organizational Intelligence as Mechanism

The claim that organizations can possess intelligence is typically treated as metaphor - a convenient shorthand for coordinated individual effort. The theoretical foundations suggest otherwise. The same phenomenon that produces individual intelligence - collective coordination through a shared medium - operates at every scale from cells to organizations. The mechanism is conserved. The substrate differs.

This section develops that mechanism through six frameworks arranged as a progressive argument: each one answers a question left open by the one before it. The first establishes what kind of system the medium is. The second explains why collective coordination through such a system produces intelligence at all. The third formalizes the mechanism mathematically. The fourth provides empirical demonstration. The fifth resolves a problem the first four leave open - how social systems achieve operational closure without physical boundaries. The sixth establishes the preconditions without which the entire mechanism fails to start. Together they specify a set of conditions that no single framework could derive alone.

A note on epistemic status. These six frameworks do not all contribute in the same way. Friston's free energy principle provides formal mechanism - mathematical structure for how self-organizing systems maintain themselves through prediction and action. Levin's scale-free cognition provides experimentally demonstrated mechanism - measurable bioelectric coordination producing collective intelligence across scales. Hutchins's distributed cognition provides empirical demonstration - observed system-level cognition in operational settings. Cohen and Levinthal's absorptive capacity provides empirically validated preconditions - measured relationships between prior knowledge and learning capacity. Maturana and Varela's autopoiesis and Luhmann's social systems theory provide structural characterization - identifying the kind of system the medium is (self-producing, communicatively closed) and the properties it must exhibit. Prigogine's dissipative structures provide a design constraint derived from thermodynamics - the requirement for continuous flow. Where a framework provides formal or experimental mechanism, we

make direct theoretical claims. Where it provides structural characterization or design constraint, we use it to specify what the medium must satisfy without claiming the biological or physical formalism transfers literally. The claim is architectural convergence, not unified theory.

4.1 Maturana and Varela: What Kind of System Is the Medium?

Maturana and Varela (1980) defined autopoiesis as self-production through operation. An autopoietic system produces the components that constitute it. A cell produces the proteins that make up the cell. The system does not exist apart from its components, nor do the components exist independently of the whole. There is a bootstrap structure: the system produces the conditions for its own continuation through its own operation.

What kind of system is the medium? It is autopoietic-like. We use this language deliberately. The translation of autopoiesis from biological to organizational systems is non-trivial and has generated sustained debate. Varela (1979) formalized the mathematical structure of autopoiesis in ways that emphasize its dependence on physical boundary conditions. Maturana himself objected to social applications, arguing that social systems lack the physical boundary that defines biological autopoiesis (Maturana, 2002). Mingers (1995) provided the most systematic analysis of the debate, concluding that social autopoiesis is defensible if the self-producing operations are identified as communications rather than physical processes — precisely Luhmann’s move. Our position draws on Luhmann’s resolution (see Section 4.5): the boundary is communicative, not physical. The medium creates conditions that approach operational closure through informational boundaries rather than physical ones.

The medium shares the core autopoietic property: it produces the conditions for its own operation. The medium generates knowledge, which is validated and enriched through operation, and in turn improves the medium’s own functioning. Remove the medium, and the capacity it supported dissipates - the organization loses cognitive properties that were dependent on the infrastructure, not on any individual component. Every organization already has a medium of sorts - the informal networks, mentoring relationships, institutional memory carried in people’s heads, the slow-motion knowledge transfer that the consulting pyramid exemplifies (Section 3). What we propose is not building a medium from nothing but extending and formalizing what already exists: making it persistent where it was ephemeral, available to computational agents where it was confined to human interaction, and operating at organizational scale where it was limited to whoever happened to be in the room. The medium is the system through which the organization maintains its cognitive identity. The design choice is whether that system remains implicit and fragile or becomes explicit infrastructure.

The implication is that the medium cannot be bolted onto existing systems as an addition. It must become the operating principle through which the organization generates and sustains its knowledge. An organization can deploy sophisticated analytical tools on top of a traditional structure and remain in Powerful Tools. The medium requires that knowledge generation and the generation of capacity for future knowledge creation are fused into a single, self-sustaining operation.

A concrete example will thread through the remaining sections, illustrating what each framework reveals about the same organizational phenomenon. Consider a mid-sized European insurer operating across three jurisdictions — property, casualty, and specialty lines. The insurer has deployed AI across claims processing, underwriting assistance, and fraud detection. Each system works competently in isolation. None perceives the organization as a whole.

Load-bearing role in the causal model: Autopoiesis specifies what kind of system the medium must be - self-producing, operationally closed, maintaining identity through its own operations. Without this, the medium is an add-on tool rather than an infrastructure that sustains itself. This determines the design requirement: the medium must generate the conditions for its own continuation, or it will decay into another static repository. For the insurer, the question is whether its AI systems form an autopoietic infrastructure — one that produces the knowledge it needs through its own operation — or a collection of tools that require manual feeding.

Autopoiesis tells us what kind of system the medium is. What it leaves unexplained is why self-production through coordination produces intelligence in the first place.

4.2 Levin: Why Does Collective Intelligence Work?

Levin's research on scale-free cognition provides the mechanism that autopoiesis describes but does not explain. Intelligence is not a property of brains. It is scale-free - emerging from collective coordination at every level from molecular networks to organisms to swarms (Levin, 2019, 2022). Each human being consists of billions of cells coordinating through bioelectric networks to produce a unified cognitive agent with goals, preferences, and memories that belong to the whole and not to its parts. There is no non-collective baseline. The unified self we experience is the emergent property of cellular coordination through shared electrical signalling.

Levin calls the coordination substrate the cognitive glue - the medium through which components coordinate. The xenobot experiments (Kriegman et al., 2020) demonstrate the principle: frog skin cells, freed from their normal bioelectric context, self-organize into entirely novel functional forms. No genetic modification - just different medium. Same cells, different coordination infrastructure, radically different collective behavior. Intelligence correlates not with the sophistication of individual components but with the density and quality of connections between them.

We should be precise about the limits of this parallel. Cells share bioelectric networks - a physical communication substrate with measurable voltage gradients. The organizational equivalent is not a physical substrate but an informational one: the knowledge infrastructure through which organizational members and agents share context, patterns, and signals. The analogy holds at the functional level - shared medium enables coordination toward collective goals - but not at the substrate level. The mechanism of coordination through shared medium is conserved across scales; the physical implementation differs.

The strategic implication is direct: investment in individual AI capability - better models, more sophisti-

cated algorithms - yields lower returns than investment in the medium that enables those capabilities to coordinate. Capability increases linearly with investment. The quality of coordination increases nonlinearly with the quality of the medium.

For the insurer: upgrading the claims AI from GPT-4 to GPT-5 improves claims summaries incrementally. Connecting the claims AI to the underwriting model through a shared knowledge infrastructure — so the claims agent perceives how underwriting decisions produced the risk profile that led to this claim — changes what the system can see. The intelligence gain comes not from better components but from richer coordination. The insurer's cognitive glue is the shared context through which claims, underwriting, and fraud detection agents perceive through each other's accumulated knowledge.

Load-bearing role in the causal model: Levin explains why investing in the coordination medium produces higher returns than investing in individual component capability. This is the theoretical basis for the fertility-first design principle: the quality of the cognitive glue determines the intelligence of the collective, not the sophistication of the parts. Paper B (Gelas and Reichhart, 2026b) applies this principle directly to the Player layer's design: agents are specified not as independent actors but as a coordinated collective whose intelligence is a property of their shared coordination infrastructure - the Map, the Physics, and their accumulated shared memory - rather than of any individual agent's capability. Remove Levin, and the argument for medium-over-capability loses its mechanistic grounding - and the Player layer collapses into a collection of independent tools rather than a collective intelligence system.

Levin explains why collective coordination produces intelligence. The missing piece is a formal account of the mechanism - how self-producing systems actually maintain themselves through coordination. Friston provides it.

4.3 Friston: How Can This Be Formalized?

Karl Friston's free energy principle provides mathematical formalization for the intuition that autopoiesis identifies and scale-free cognition explains. Any system that persists in a changing environment must continuously update its internal model to match reality - minimizing the divergence between expectations and observations (Friston, 2010). The system that fails to update is surprised by the environment, and persistent surprise is another word for disintegration. This is autopoiesis rendered in probability theory: the self-producing system maintains itself by keeping its internal model aligned with its environment.

Friston provides formal tools the other frameworks lack. Markov blankets - statistical boundaries separating inside from outside - give a precise definition of operational closure (Kirchhoff et al., 2018). Where Maturana and Varela describe operational closure as a property and Levin shows how it scales, Friston specifies its mathematical structure. The Markov blanket is the interface through which the system senses and acts. For an organism, sensory and motor systems form the blanket. For an organization, the medium forms it - creating informational boundaries that define what is inside organizational cognition versus what is external environment.

Active inference extends the picture in a way that connects back to Levin’s cognitive glue. Self-maintaining systems do not merely update their models passively. They act on the environment to make observations less surprising. An organism explores, reducing uncertainty. An organization experiments, tests, gathers intelligence. The medium enables organizational active inference - the infrastructure for continuous exploration, hypothesis testing, and model updating. Without the medium, active inference is episodic. With it, active inference becomes the continuous cognitive coordination that Levin’s framework predicts.

The insurer’s Markov blanket is defined by what its medium can sense and act upon. Currently, the blanket is thin: the claims system senses claims, the underwriting system senses applications, the fraud system senses anomalies. Each has its own narrow boundary with the environment. With the medium, the blanket thickens: the organization senses through the combined perception of all three systems, connected through shared context. Active inference becomes organizational — the insurer continuously updates its model of risk, not through periodic actuarial review but through the real-time accumulation of cross-system patterns. A rising claims frequency in a specific postcode, combined with underwriting data showing a new contractor operating there, combined with fraud patterns flagging that contractor’s work quality — the medium makes this convergence perceivable. No individual system sees it. The organization’s Markov blanket, defined by the medium, does.

We are explicit about limits. Friston’s framework was developed for biological systems with continuous-time dynamics and physical embodiment. Its application to organizational systems is an extension, not a derivation. Ramstead, Badcock, and Friston (2018) have begun this extension formally, proposing that the free energy principle applies to social systems through shared generative models — groups that coordinate effectively do so by maintaining aligned predictions about each other’s behavior. Veissière et al. (2020) develop this further through “thinking through other minds,” arguing that cultural practices function as shared priors that reduce collective uncertainty. These extensions support our application but do not settle it. Whether organizational systems literally minimize variational free energy in Friston’s technical sense is an empirical question we do not claim to have answered.

Load-bearing role in the causal model: Friston provides the formal mechanism connecting perception, prediction, and action that the previous two frameworks describe qualitatively. The Markov blanket defines the medium’s boundary - what is inside organizational cognition and what is external. Active inference explains how the medium learns: not passive storage but continuous prediction and correction. Remove Friston, and the medium has no account of how it maintains itself through operation. The framework also opens a natural extension: if each organization defines its Markov blanket through its medium, then inter-organizational intelligence emerges from the interaction between blankets - organizations sensing and acting on each other through governed interfaces. Paper B (Gelas and Reichhart, 2026b) develops this as the inter-organizational architecture, where counterparties, regulators, and market infrastructure interact as coupled world models, each maintaining its own tripartite structure while exchanging state information through defined boundaries.

Three frameworks now specify the mechanism: a self-producing system that generates intelligence through collective coordination, maintaining itself through predictive modeling and active inference. But this remains theoretical architecture. Has anything like it been observed in practice?

4.4 Hutchins: Has This Been Demonstrated?

Edwin Hutchins (1995a) provided the empirical demonstration decades before the theoretical frameworks converged. Studying navigation teams aboard naval vessels, Hutchins (1995) showed that cognitive systems have properties not reducible to individual cognition. No individual crew member knows the ship's position. The navigator does not know it. The radio operator does not know it. The lookouts do not know it. But the system - the coordinated operation of navigator, radio operator, lookouts, and instruments - computes it with precision.

The cognitive work distributes across people and artifacts. Representations transform as they propagate through the system: bearing sighted, position converted, plotted on chart, combined with other fixes, course recommended. Hutchins's observation about aircraft cockpits crystallizes the point: speed bugs do not help pilots remember speeds; they are part of the process by which the cockpit system remembers speeds. The artifact participates in cognition - it is part of what thinks.

This is the three preceding frameworks made visible in a single empirical case. Levin's cognitive glue: the navigation system's intelligence is a property of the coordination, not of the individuals. Friston's active inference: the system continuously updates its position model by sensing (taking bearings), predicting (dead reckoning), and correcting (comparing prediction to observation). Autopoietic self-production: the system produces the navigational knowledge that justifies and sustains the system itself.

The insurer's equivalent: no individual system knows the organization's risk position. The claims system does not know it. The underwriting system does not know it. The actuarial model does not know it. But the medium — claims data cross-referenced with underwriting decisions, enriched by fraud detection patterns, grounded in regulatory obligations — computes a risk position more accurate than any individual system's view. The knowledge lives in the relationships between systems, not in any system. Like Hutchins's navigation team, the insurer becomes a cognitive system — one that knows what none of its components know.

Load-bearing role in the causal model: Hutchins provides empirical proof of concept. Without this framework, the entire argument is theoretical - plausible but undemonstrated. Hutchins shows that system-level cognition is observable, measurable, and real. The navigation team is what the medium looks like when it works: knowledge that lives in the system, not in any individual head.

Hutchins's demonstration is powerful but bounded. Hollan, Hutchins, and Kirsh (2000) extended the framework beyond co-located teams, arguing that distributed cognition applies wherever representational states propagate across media and agents — including computational ones. Zhang and Norman (1994) formalized the role of external representations in cognitive systems, showing that the distribution of infor-

mation between internal and external representations fundamentally changes the nature of the cognitive task. These extensions matter because the medium must operate at organizational scale: across continents, asynchronously, and with computational as well as human agents. The mechanism needs to work when boundaries are informational rather than physical.

4.5 Luhmann: How Do Social Systems Achieve Closure Without Physical Boundaries?

Niklas Luhmann (1995) resolves the problem that autopoiesis leaves open and Hutchins's empirical example does not address: how can social systems be operationally closed when they lack the physical boundaries that define biological autopoiesis?

For Luhmann, social systems are constituted by communications, not by people. A social system maintains itself by producing communications from communications, establishing a boundary between system and environment through its own communicative operations. The environment does not directly cause the organization to change. It generates signals. The organization responds to those signals as processed through its own communicative structure. Two organizations in identical external conditions will respond differently because their internal communication structures differ.

This fills the gap. Maturana and Varela require operational closure but define it through physical boundaries. Luhmann shows that communicative closure achieves the same function: the system reproduces itself through its own operations, but the operations are communications rather than biochemical processes. The medium is the infrastructure that creates this communicative operational closure - both the Markov blanket through which the organization senses and acts (Friston) and the communicative system through which the organization maintains itself (Luhmann). Two descriptions of the same mechanism viewed from different theoretical angles.

Load-bearing role in the causal model: Luhmann resolves the boundary problem that would otherwise break the analogy to biological systems. Without communicative closure, the medium has no account of how organizational identity is maintained - and the autopoietic characterization collapses into metaphor. Luhmann makes it mechanism: the medium is the communicative infrastructure through which the organization reproduces itself.

All five frameworks share a blind spot. They assume a system already in operation - already coordinating, already maintaining its Markov blanket, already communicating. None addresses how the system reaches that state. The mechanism requires prior capacity, and that capacity has to come from somewhere.

4.6 Cohen and Levinthal: What Must Exist Before the Mechanism Can Start?

Cohen and Levinthal (1990) introduced absorptive capacity as the organizational ability to recognize, assimilate, and apply new external knowledge. Their research established that this ability depends heavily on prior related knowledge. An organization rich in domain knowledge can recognize and integrate new knowledge rapidly. An organization lacking that prior base will struggle regardless of how capable its

tools are.

Absorptive capacity is cumulative and path-dependent. Prior knowledge structures determine what new knowledge the organization can even perceive - let alone integrate. This is the precondition the other five frameworks assume but do not specify. Autopoiesis requires components that can produce themselves. Scale-free cognition requires coordination mechanisms. Active inference requires a model to update. Distributed cognition requires representational pathways. Communicative closure requires communicative infrastructure. Cohen and Levinthal identify what determines whether any of these conditions can be met: the organization's accumulated prior knowledge.

Load-bearing role in the causal model: Cohen and Levinthal specify the bootstrap condition - the precondition without which none of the other five mechanisms can start. Without absorptive capacity, the medium has nothing to work with: no prior knowledge to coordinate, no patterns to propagate, no base from which active inference can begin. This is what connects fertility (Section 2) to the mechanism stack: fertility is the measure of whether the bootstrap condition is met.

Fertility is absorptive capacity operationalized. The six dimensions described in Section 2 - knowledge flow rate, contextual richness, pattern diversity, temporal depth, integration maturity, governance completeness - are the specific architectural features that determine whether an organization can bootstrap the autopoietic, scale-free, actively inferring, distributed, communicatively closed cognitive system the other five frameworks describe. You cannot produce organizational intelligence by adding capable agents to an organization that lacks the prior knowledge infrastructure. You must build the absorptive capacity that enables the medium to bootstrap itself into self-sustaining operation. Once that threshold is crossed, capability investments compound. Before it, they do not.

The six frameworks form a single integrated claim that none supports alone. The medium is an autopoietic-like system that produces organizational intelligence through the same scale-free coordination mechanism Levin identified in biological systems - operating here through informational substrates rather than bio-electric networks. Friston's free energy principle formalizes how such a system maintains itself: active inference through an organizational Markov blanket. Hutchins demonstrated empirically that distributed cognitive systems exhibit exactly these properties. Luhmann resolved the boundary problem: communicative closure substitutes for the physical boundaries biology requires. Cohen and Levinthal specified the bootstrap condition: prior accumulated knowledge determines whether the mechanism can start at all. Remove any one condition, and the mechanism fails in a specific, identifiable way. The medium is the infrastructure that satisfies all six simultaneously.

A natural objection: ten frameworks is too many. The synthesis could be theoretical accumulation rather than structural necessity. The test is subtraction. Drop Friston, and the system has no formal account of how it maintains itself - self-production and coordination are described qualitatively but without the mechanism connecting perception, prediction, and action. Drop Hutchins, and the entire argument is theoretical with no empirical demonstration that distributed cognition actually produces system-level intel-

ligence. Drop McGilchrist, and the framework implies eventual full automation - missing the categorical boundary that makes human participation permanent. Each framework was included because removing it leaves a specific gap the remaining nine cannot fill.

What remains to be specified is the relationship between this mechanism and what the AI industry calls maturity - and why the industry's dominant measure - autonomy - captures only part of what organizational intelligence requires.

5. From Autonomy to Initiative

Two AI agents operate inside the same insurance company. Same model, same capabilities, same access to data. The first monitors claims for fraud. It is excellent - flagging anomalies, cross-referencing patterns, escalating suspicious cases with a precision no human team could match at that volume. It does exactly what it was told to do. The second agent notices that a cluster of legitimate claims from a specific region share characteristics suggesting a supplier quality problem nobody asked about. It flags not fraud but an emerging product liability risk - three months before the first lawsuit arrives.

Same capability. Different relationship to goals. The first agent is autonomous: it executes assigned objectives independently. The second has initiative: it identifies objectives worth pursuing. The gap between them is contextual. The second agent perceives through organizational knowledge - claims history, supplier relationships, regulatory patterns, the company's risk appetite - what matters before anyone specifies it.

The AI industry defines maturity as increasing autonomy. Level 5 in most frameworks means pervasive autonomous capability - AI making decisions without human involvement. This gets the direction wrong. Autonomy means independence in how to work but dependence on humans for what to work on. Give an autonomous agent a well-defined objective and it will find the optimal path. It will never notice that the objective itself has become obsolete. Initiative means perceiving what is worth doing - aligned with organizational purpose not through instruction but through immersion.

The consulting pyramid makes this distinction visceral. A junior consultant fresh from an MBA can autonomously execute any analysis you assign. Hand her a dataset and a question and she will run the regression, build the deck, present the findings with competence. She is autonomous. She is also entirely dependent on someone telling her which question to investigate. A senior partner walks into the same client meeting with the same data available and perceives something different. She notices that the CFO's question about cost reduction is actually a question about organizational restructuring the board has not yet authorized. She sees that the real problem is not the one on the engagement letter. She redirects the conversation toward what the client actually needs rather than what was formally requested.

Raw intelligence doesn't explain it - the junior may have the higher IQ. Nor does capability - they have access to the same analytical tools. What the partner has is twenty years of contextual immersion: thousands of client situations, political dynamics, industry patterns, outcomes observed over time. She has absorbed how organizations actually work - not as a set of rules she can recite but as a felt sense for what matters, what to probe, when the stated problem masks the real one. That felt sense is what produces initiative. And it developed through immersion, not through anyone handing her a more detailed instruction manual. The medium compresses this from decades to months by providing the same immersive context - accumulated patterns of how the organization operates, what it values, what outcomes it seeks - now available to computational agents.

The obvious counterargument: why not just give agents more goals? Expand the objective function. Instead of "flag fraud," specify "flag fraud, supplier risks, regulatory exposure, and emerging liability patterns." Add more instructions. Write better prompts.

This fails, for a precise reason. The insurance agent that spotted the supplier quality problem did not operate from a more detailed goal specification. It perceived through organizational context that supplier quality problems matter to this specific company - because of its claims history, its risk tolerance, its regulatory position, its competitive strategy. A different insurer with different context might correctly ignore the same signal. Initiative is contextual perception. You cannot specify in advance what an agent should notice, because the value of noticing something depends on organizational context that shifts continuously. The space of things worth noticing is effectively infinite. No instruction set can enumerate it. This is why the transition from autonomy to initiative requires immersion, not instruction - the same reason no consulting firm has ever produced a manual that turns juniors into partners.

The mechanism already exists in AI development, though it has not been applied at organizational level. The progression from pre-training through reinforcement learning from human feedback (RLHF; Christiano et al., 2017) to constitutional AI (Bai et al., 2022) maps directly: pre-training builds base capability, supervised fine-tuning teaches procedures, RLHF develops preferences through feedback, constitutional AI internalizes values that guide action. Most organizations are stuck at stage two with their agents - explicit instructions, defined workflows, rules-based governance - while expecting stage four outcomes. The medium is organizational RLHF: accumulated patterns of what the organization values, enabling agents to learn judgment rather than follow procedures.

Rules-based governance - specifying what agents can and cannot do - works for autonomy but cannot produce initiative. More capability means more rules, and rules substitute for perception rather than enabling it. Context-based governance is the alternative: the agent perceives what matters because it is immersed in organizational patterns. Boundaries become implicit in how the organization makes trade-offs, what it will and will not do. This is how senior professionals are governed in practice - not through rulebooks but through contextual immersion until they perceive through organizational eyes.

The strategic implication cuts to the core of competitive advantage. Capability commoditizes - foundation

models converge in capability faster than the market expects, agent frameworks follow the same trajectory, and autonomous execution is on course to become table stakes within a few years. What will not be available off the shelf is the organizational context that produces initiative. Your accumulated patterns, tacit knowledge traces, decision histories, and ways of operating instantiate your specific organizational identity. Agents immersed in your medium develop initiative aligned with your organization's nature. That alignment cannot be purchased from a vendor. It cannot be replicated by a competitor running the same models on the same cloud infrastructure. It can only be grown through the accumulation of organizational patterns over time. Initiative requires a medium, and the medium is - by definition - unique to the organization that built it. This is the one form of AI advantage that does not erode with the next model release. But initiative requires a medium that stays alive - one that continuously generates and refreshes the patterns through which agents develop contextual perception. A dead medium produces no initiative, regardless of how rich it once was.

6. The Medium as Living System

What keeps the medium alive? Flow. The medium is a dissipative structure: it maintains itself through continuous energy flow, not through preservation. Stop the flow, and the structure collapses - regardless of how much knowledge it contains. This is the design constraint that separates a medium capable of sustaining initiative from another failed knowledge repository.

Ilya Prigogine's Nobel Prize-winning work established that open systems far from thermodynamic equilibrium can spontaneously generate order - not despite energy flow but through it (Prigogine, 1980; Prigogine and Stengers, 1984). Dissipative structures maintain their organization by continuously importing energy and exporting entropy. The requirements are specific: the system must be open, operating far from equilibrium, governed by nonlinear dynamics, and maintained by continuous flow. Equilibrium structures like crystals or archived files maintain order in isolation. Dissipative structures like organisms - or the medium - maintain order through flow. Stop the flow, and the structure collapses.

This provides the thermodynamic explanation for why knowledge management has persistently underperformed. Most knowledge management systems create closed systems - capturing knowledge then isolating it in repositories. These systems tend toward maximum entropy by the second law. Archived knowledge decays toward uselessness not through any specific failure but through the fundamental physics of isolated systems. The failure rates of 50-70% are not management failures. They are thermodynamic inevitabilities. Traditional KM builds equilibrium structures in a far-from-equilibrium context. They decay by necessity.

The medium must remain open and continuously fed. Quality input must exceed entropy generated. Curation is essential - the medium needs negentropy sources. The computational and organizational cost

is real, but returns compound as the medium matures.

The knowledge flywheel illustrates the mechanism. In the first turn, AI systems extract and structure patterns from organizational data and operations. In the second turn, human teams curate, validate, and enrich these patterns with contextual judgment the system cannot supply. In the third turn, this enriched knowledge becomes the operating environment for AI agents. In the fourth turn, better-contextualized agents discover richer patterns. The cycle repeats, and each revolution enters the spiral transformed.

This flywheel is a dissipative structure: it maintains order through continuous energy flow. It degrades without input. It generates increasing order - richer, more connected, more contextually dense knowledge - through operation. It exhibits the nonlinear dynamics characteristic of dissipative structures: small improvements in knowledge quality produce disproportionate improvements in agent performance, which accelerate knowledge capture, which further improves quality. The medium is the infrastructure that enables and sustains this flywheel at organizational scale.

The insurance company from Section 5 illustrates the point concretely. The agent that spotted the supplier quality problem perceived it through continuously accumulated patterns - claims cross-referenced with supplier data, regional outcomes, liability trajectories. Those patterns existed in the relationships between recent claims, relationships that only a living medium could maintain. Archive the data and deploy agents against a static repository, and the signal disappears.

The organizational implications of operating far from equilibrium are well-documented. Snowden and Boone (2007) formalized this through the Cynefin framework, distinguishing simple, complicated, complex, and chaotic domains — each requiring fundamentally different management approaches. The medium operates in the complex domain: cause and effect are discoverable only in retrospect, and the system's behavior emerges from the interactions of its components rather than from any central design. Stacey (2001) developed a parallel analysis through complex responsive processes, arguing that organizational knowledge is not a thing that can be stored but a pattern of relating that exists only in the ongoing interactions between people. Both analyses converge on the same design requirement: the medium cannot be a static structure. It must support emergence.

Prigogine's bifurcation theory adds a critical dimension. Under sufficient stress, dissipative structures do not degrade linearly. They reach bifurcation points where they either transform to a higher-order structure or disintegrate. There is no middle path at the bifurcation. The medium can play three roles at such moments: enable transformation by surfacing patterns that help the organization see through the crisis, accelerate collapse by preserving patterns that no longer match reality, or trigger bifurcation by surfacing contradictions the organization cannot ignore. Which role it plays depends on whether the medium supports pattern exploration - not just pattern preservation.

This has a direct design implication. The medium must support the discovery of new patterns, not just the efficient retrieval of established ones. It must maintain competing interpretations rather than con-

verging prematurely on organizational consensus. It must embed mechanisms for challenging what the organization believes - minority views that persist at low confidence rather than being deleted, external signals that disrupt comfortable assumptions, temporal decay that prevents stale knowledge from masquerading as current truth. Without these features, the medium becomes an accelerant toward whatever the organization already was, amplifying strengths and pathologies alike. The claims-based architecture described in Paper B (Gelas and Reichhart, 2026b) is the specific technical mechanism that addresses this requirement - knowledge as weighted assertions with provenance, temporal decay, and competing interpretations rather than as binary truth in a static repository. Paper B provides the engineering specification: how claims are represented, how prevalence weighting works, how dissent is structurally preserved, and how the dynamics layer enforces causal reasoning that pure neural approaches cannot supply. The theoretical requirement established here - that the medium must support pattern exploration, not just pattern preservation - is what drives Paper B's architectural decisions.

The thermodynamic perspective explains a persistent failure pattern. Dissipative structures require time to reach the bifurcation threshold - a pre-bifurcation phase where order accumulates beneath visible operation. During this phase, the system appears stuck. Measured against project-based timelines, it looks like a failing investment. Resources get withdrawn. The system collapses. The conclusion that the approach failed is wrong - the approach was interrupted before the dynamics it was designed to produce could manifest.

7. The Categorical Limits of Computation

A well-fed medium can grow indefinitely in richness and scale. It cannot grow past a categorical boundary that no amount of energy or sophistication will cross. Recognizing this boundary is a design requirement, not a concession of weakness.

Iain McGilchrist's hemispheric analysis makes the sharpest version of this argument. Across *The Master and His Emissary* (McGilchrist, 2009), *The Matter With Things* (McGilchrist, 2021), and his 2024 essay "Resist the Machine Apocalypse" (McGilchrist, 2024), McGilchrist demonstrates that the two cerebral hemispheres approach reality in categorically different ways - not in the folk-psychology sense of left-brain/right-brain stereotypes, but in a deeply researched account of how attention shapes what we perceive.

The left hemisphere manipulates tokens and symbols. It focuses narrow-beam attention on targets already identified. It sees parts rather than wholes, deals in the explicit, abstract, and decontextualized, creates self-consistent systems that do not know what they do not know, and prioritizes manipulation, control, and utility. AI does all of this, at scale, with extraordinary speed.

The right hemisphere does something categorically different. It maintains broader attention open to what-

ever may be. It perceives uniqueness and context-dependence. It stays in touch with experience itself rather than representations of experience. It understands the implicit, the ambiguous, and the metaphorical. It processes betweenness - meaning that exists in relationships rather than in things. It connects to embodied, lived experience. McGilchrist's claim is that what the right hemisphere does cannot be replicated through computational approaches grounded in symbolic manipulation. The right-hemisphere gap is a categorical difference in kind.

This characterization requires qualification. Pure neural networks - including large language models - are left-hemisphere processors in McGilchrist's taxonomy: pattern-matchers operating on statistical regularities, effective at token manipulation, blind to causation. But neuro-symbolic architectures — what Garcez and Lamb (2023) call “the 3rd wave” of AI — extend what computation can accomplish within the left-hemisphere domain in ways that matter. Symbolic reasoning is not pattern matching. It is logic - causal chains, constraint enforcement, if-then relationships that hold because they are structurally necessary, not because they are statistically likely. The neuro-symbolic approach described in Paper B (Gelas and Reichhart, 2026b) - neural perception combined with symbolic reasoning - extends the left hemisphere's reach from statistical prediction to logical necessity. This is precisely what regulated enterprises require: decisions whose correctness is provable rather than merely probable.

The distinction matters for medium design. The dynamics layer can handle organizational logic that pure language models cannot: regulatory constraints that must be enforced categorically, dependency chains that must be traced deterministically, compliance decisions that must be auditable to their logical premises. This does not cross McGilchrist's categorical boundary. Symbolic reasoning remains left-hemisphere - it still operates on representations rather than experience, on parts rather than wholes. But it extends the left hemisphere's reach significantly further than pure neural approaches.

The right-hemisphere gap remains categorical. No current computational approach - neural, symbolic, or hybrid - replicates the capacity to perceive uniqueness, to hold ambiguity without resolving it, to understand meaning that exists in relationships rather than representations. This is why human participation is a design requirement for the medium, not a transitional concession that will be engineered away. Humans supply what computation cannot: grounding in physical and social reality, affective feedback that signals what matters before analysis can explain why, contextual judgment that adapts principles to unique situations, and imagination - the capacity to perceive possibilities that pattern recognition alone would miss.

The insurer's through-line makes this concrete. The medium can detect the statistical convergence — rising claims, problematic contractor, geographic concentration — through left-hemisphere pattern matching. It can trace the regulatory consequences through symbolic reasoning. It can flag the anomaly, calculate the exposure, and recommend the action. What it cannot do is perceive that the real issue is not the contractor but a shifting relationship between climate-driven subsidence and the insurer's historical pricing assumptions in that region — an insight that requires holding the ambiguity between multiple ex-

planatory frames simultaneously, sensing the felt significance of a pattern that does not yet have a name. The underwriting director who has spent twenty years in that market perceives this. The medium cannot. The medium can, however, surface the patterns that make the underwriting director's perception actionable at organizational scale — connecting her insight to claims data, actuarial models, and reinsurance thresholds she could not traverse alone. Neither alone reaches the conclusion. Together, they do.

McGilchrist goes further. As machines become more like humans, humans risk becoming more like machines - shaped by the obligation to interact with mechanical systems that reward left-hemisphere behavior. If the medium captures only what can be processed computationally, it gradually reshapes organizational culture in its own image. Organizations may come to trust patterns over intuition, value documentation over craft, prefer explicit over tacit, choose mediated knowledge over direct experience. The medium must function as emissary, not master - McGilchrist's formulation of the proper relationship between hemispheres. The moment what the medium says overrides human judgment categorically, the system has failed its own design principles.

This argument connects to Kauffman's adjacent possible in a way that clarifies the relationship between computation and creativity. Kauffman (2019) defines the adjacent possible as the set of outcomes reachable from the current state through recombination of what already exists. The concept originates in chemistry - the set of molecular species one reaction step away from those already present - but extends to any system where current configuration determines what configurations are next available.

Kauffman illustrates the point with the screwdriver argument: there is no algorithm that can list all the uses of a screwdriver, because the uses depend on the diversity of the context. A screwdriver in empty space has few uses. The same screwdriver in a full garage has many. Capability is the sharpness of the screwdriver. Fertility is how full the garage is. Same capability, vastly different possibilities depending on context richness. AI excels at exploring the adjacent possible through recombination - it can identify combinations of existing patterns at speeds and scales no human team could match. What AI struggles with is perceiving truly new possibilities - the kind that emerge not from recombination of the known but from the imaginative leap that sees what has never been. The distinction maps onto McGilchrist's hemispheric analysis precisely: recombination is left-hemisphere work (algorithmic, exhaustive, pattern-based), while imagination is right-hemisphere work (integrative, contextual, perceptual).

Human imagination is what enables the organization to perceive opportunities in the adjacent possible that pattern recognition alone would miss. The medium expands the adjacent possible by increasing the diversity and richness of context available for recombination. Human imagination perceives what that expanded possibility space contains - the novel configurations that no algorithm would generate. Neither alone produces the highest-order organizational outcomes. Together, they enable the organization to generate value it could not have specified in advance. This is why the medium must be designed as a human-AI collaborative system, not as an automation layer that progressively reduces human involvement.

8. The Cost Curve of Waiting

Path dependency creates strategic consequences with compounding effects. Each capability optimization that does not account for fertility hardens the architecture, and each hardening raises the cost of adding fertility later.

Four mechanisms compound. First, workflow entrenchment: processes designed for autonomous execution become organizational defaults that resist redesign. The fact that 42% of companies scrapped most of their AI initiatives in 2025 (S&P Global, 2025) suggests that even early investments create enough organizational commitment that abandonment, not evolution, becomes the response to failure. Second, integration fabric hardening: technical architectures optimized for speed - low-latency agent execution, high-throughput data pipelines - resist reconfiguration toward the richer, relationship-preserving structures that fertility requires. Third, talent orientation: teams trained for agent deployment and operational excellence develop competencies and incentive structures misaligned with exploratory, fertility-first approaches. Fourth, data architecture lock-in: organizations that structure knowledge as documents for retrieval face substantial migration costs to claims-based architectures. The move from document-based to graph-based to claims-based knowledge representation is a progression of increasing structural richness, and each transition requires not just technical migration but organizational reconceptualization of what knowledge is.

The insurance company from Section 5 illustrates all four mechanisms at once. Claims workflows optimized for processing speed resist redesign toward contextual pattern capture. The data architecture treats each claim as an independent event rather than a node in a relationship network. The talent base is trained for claims adjudication, not knowledge curation. The longer these defaults operate, the harder it becomes to add the contextual richness that produced the supplier quality insight.

The cost curve is not linear. It accelerates. Sydow, Schreyögg, and Koch (2009) describe the mechanism precisely: self-reinforcing dynamics in path dependence do not merely add costs - they multiply them. Each mechanism reinforces the others. Entrenched workflows require specific integration patterns, which require specific talent, which generates specific data architectures, which further entrench the workflows. The compounding is multiplicative, not additive.

The compounding creates a visibility problem. Capability investment generates returns that are visible, immediate, and easily attributed - a chatbot answers questions faster, an agent processes claims more efficiently, a model generates code at lower cost. Fertility investment generates returns that are diffuse, delayed, and difficult to attribute - the organization discovers opportunities it would not otherwise have seen, avoids risks it did not know it was running, generates strategies no one specified. This visibility asymmetry biases investment toward capability even when the theoretical and empirical case favours

fertility. Each visible capability win reinforces the investment pattern that produced it, while the invisible fertility deficit grows beneath the surface - deepening the very path dependency that makes later correction expensive.

Capability will arrive whether organizations build it or not - vendors will deliver it, costs will fall. Fertility will not arrive on its own. It must be deliberately cultivated, and the window for cultivation narrows as capability-first architecture hardens. Those that wait face a retrofit: possible, but expensive, slow, and competing against organizations already generating value they could not have specified in advance. Every organization will eventually need the medium. Whether the cost will still be bearable when the need becomes undeniable is the strategic gamble.

9. Discussion

9.1 Theoretical Contribution

The paper's central contribution is the demonstration that ten independently established theoretical frameworks, when synthesized, specify a mutually reinforcing set of conditions for organizational intelligence. Each framework answers a question the others leave open:

Framework	Question Answered
Maturana and Varela	What kind of system is the medium?
Levin	Why does collective intelligence work?
Friston	How does the system maintain itself?
Hutchins	Has distributed cognition been demonstrated?
Luhmann	How do social systems sustain operational closure?
Cohen and Levinthal	What are the preconditions for learning?
Nonaka and Takeuchi	How does knowledge convert between forms?
Kauffman	What can emerge from a fertile system?
Prigogine	What does the system need to survive?
McGilchrist	What are the categorical limits?

The convergence is structural, not decorative. Remove any single framework and the mechanism fails in a specific way: without Luhmann, operational closure has no account for non-physical boundaries and the medium cannot define organizational identity. Without Prigogine, there is no explanation for why knowledge repositories decay and no design constraint for keeping the medium alive. Without Cohen and Levinthal, there are no preconditions - the mechanism has no bootstrapping account. Each framework is load-bearing. The contribution is not that these frameworks exist individually but that their

combination specifies architectural requirements no single framework could derive - requirements that current AI maturity frameworks do not address.

The autonomy-to-initiative distinction represents a second theoretical contribution with direct practical consequences. Current AI maturity models define their highest level as pervasive autonomous capability - agents executing without human involvement. This paper redefines Level 5 as agents that identify goals worth pursuing through immersion in organizational context. The redefinition changes what maturity means: the measure shifts from independence of execution to quality of perception. It changes what organizations should invest in: not more capable agents but richer context. And it changes what success looks like: not faster execution of assigned goals but discovery of goals worth assigning. Current evaluation benchmarks do not measure initiative. Developing instruments that distinguish autonomous execution from contextual goal identification is an immediate research priority.

9.2 Scope and Limitations

As a theory paper, this work constructs a framework rather than testing it. Several scope boundaries are important.

The biological grounding uses “autopoietic-like” to describe the medium’s properties. The claim is specific: the medium creates conditions that exhibit operational closure through informational boundaries rather than physical ones. This is consistent with Luhmann’s extension but more cautious than claiming full biological autopoiesis.

The path dependency claim describes a cost curve, not a wall. Organizations can always restructure. The claim is that switching costs compound - making the retrofit progressively more expensive, not impossible.

The relationship between the medium and organizational intelligence is enablement, not causation. The medium creates conditions where organizational intelligence can emerge. It does not produce it automatically. Whether emergence occurs depends on the quality of the medium, the capability of the agents, the quality of human participation, and the organization’s existing absorptive capacity. Kauffman’s concept of enablement is the appropriate framing: the medium enables new possibilities without causing specific outcomes.

The Polanyi claim requires careful statement. The medium dramatically expands the boundary of what can be captured from organizational tacit knowledge through behavioral observation and pattern extraction. It does not dissolve Polanyi’s paradox. A residual tacit dimension remains - and the medium must be designed to acknowledge and point toward that dimension rather than pretending to have captured it.

The authors are currently applying elements of the architecture described in Paper B within regulated financial services engagements. Empirical results will be reported in forthcoming papers, providing the first opportunity to test several of the propositions generated by this theoretical framework — particularly

the Fertility Index, the knowledge flywheel effect, and the autonomy-to-initiative transition.

A further limitation deserves explicit treatment: the amplification risk. The medium amplifies what is already there - including what should not be amplified. If the medium captures how an organization thinks, and that thinking embeds bias, blind spots, or outdated assumptions, the medium does not merely preserve those flaws. It makes them structural. Amazon's hiring algorithm, trained on historical patterns, penalized the word "woman" and favoured male applicants (Dastin, 2018). The COMPAS recidivism system falsely flagged Black defendants as high risk more often than white defendants (Angwin et al., 2016). These are cases where pattern amplification reproduced and scaled existing bias. The medium-specific risk is that pathologies operate at the level of tacit pattern propagation, not just explicit content. The medium does not merely inherit bias in documents - it can inherit bias in the relationships between claims, in what gets weighted highly and what decays. Edmondson's (1999) research on psychological safety is directly relevant: knowledge sharing requires conditions of trust that cannot be engineered through incentive design alone. Without psychological safety, the medium captures only what people feel safe contributing - which may systematically exclude the dissenting perspectives and minority views the organization most needs. The claims-based architecture described in Paper B addresses this through structural mechanisms - minimum prevalence floors, dissent mechanisms, external pattern injection - but the social conditions for healthy contribution remain an organizational challenge, not a technical one.

9.3 Research Agenda

The framework generates several testable propositions that constitute the research agenda:

First, empirical validation of the Fertility Index. The six dimensions - knowledge flow rate, contextual richness, pattern diversity, temporal depth, integration maturity, governance completeness - require operationalization and measurement. The prediction is that organizations scoring higher on fertility will show higher AI ROI than organizations scoring higher on capability, controlling for total investment.

Second, longitudinal studies of organizations building medium-like infrastructure. The prediction is that these organizations will demonstrate the characteristic dynamics of dissipative structures: an apparent plateau followed by rapid transformation at the bifurcation point. Cross-sectional studies will miss this dynamic because the pre-bifurcation phase appears indistinguishable from failure.

Third, formal modeling of the claims architecture described in Paper B. The prediction is that claims-based knowledge architectures with prevalence weighting will maintain greater pattern diversity over time than document-based or retrieval-augmented alternatives - and that this diversity will correlate with higher rates of novel value generation.

Fourth, cross-industry comparison. Regulated industries, where compliance requirements inadvertently build much of the infrastructure the medium requires, should show a shorter path from compliance infrastructure to organizational intelligence than unregulated industries starting from greenfield. The prediction is that regulatory pressure is, counterintuitively, a fertility driver.

Fifth, the autonomy-to-initiative transition. The prediction is that agent performance measured by goal identification - not just goal execution - will correlate with the richness of organizational context available to the agent, independent of the agent's raw capability. This requires developing measurement instruments that distinguish autonomous execution (correctly completing assigned tasks) from initiative (identifying which tasks are worth completing) - a distinction that current AI evaluation benchmarks do not make.

Sixth, the knowledge flywheel effect. The prediction is that self-reinforcing improvement in knowledge quality and agent performance will be measurable in organizations with medium-like infrastructure and absent in those with static knowledge repositories. Longitudinal measurement would need to capture whether each cycle of operation measurably enriches the medium or whether the system reaches a plateau, and if so, under what conditions bifurcation to a higher-order state occurs versus degradation.

9.4 Relationship to Companion Papers

This paper provides the theoretical foundations. Paper A (Reichhart and Gelas, 2026a) diagnoses the specific failure mechanism - dynamics blindness - that explains why the vast majority of enterprise GenAI pilots fail to meet expectations. Dynamics blindness is the inability of current AI to reason causally about how organizational rules interact with dynamic systems. The theoretical framework developed here explains why dynamics blindness is a structural property of the current paradigm, not a capability shortfall awaiting better models.

Paper B (Gelas and Reichhart, 2026b) specifies the architecture - a tripartite structure of state (Map), dynamics (Physics), and agency (Player) operating on a claims-based knowledge representation with prevalence weighting. The theoretical framework developed here explains why that specific architecture is required: it is the minimal structure that satisfies all ten theoretical conditions simultaneously. The Map provides the substrate for distributed cognition. The Physics enables active inference and causal reasoning - consuming state sequences and observations to trace both normative consequences and predicted trajectories. The Player layer is designed as a coordinated collective - individual agents, specialized teams, and swarms that develop organizational perception through episodic memory, long-term context accumulation, and structured self-critique - where autonomy operates and initiative can eventually emerge. Paper B further extends the architecture to inter-organizational boundaries, applying Friston's Markov blanket framework to specify how organizations interact as coupled world models. The claims architecture maintains the medium as a dissipative structure with competing interpretations rather than a static repository tending toward entropy.

Together, the five papers form a complete programme: diagnosis (A), architecture (B), theoretical foundations (C), practitioner methodology (D), and the capstone argument for initiative as the ultimate measure of AI maturity (E). Each paper is designed to stand alone while gaining additional force from the others. Paper A can be read as a diagnostic of current enterprise AI failures without reference to the theoretical apparatus developed here. Paper B can be read as an architectural specification without the biological and cognitive grounding. This paper can be read as a contribution to organizational theory without the

specific enterprise application. Paper D can be read as a governance methodology without the theoretical underpinning. Paper E can be read as a standalone argument about the autonomy-to-initiative transition. The combined programme produces something none achieves individually: a theoretically grounded, architecturally specified, empirically motivated account of what organizational intelligence requires in the age of AI.

10. Conclusion

Every organization already operates through a medium. The question this paper answers is what that medium must become - and what ten theoretical traditions say about why.

The answer is architectural. Ten frameworks converge on a set of engineering requirements: the medium must be self-producing, scale-free in its coordination, continuously predictive, distributed across human and computational agents, communicatively closed, grounded in prior knowledge, capable of converting between tacit and explicit forms, expansive of possibilities, sustained through energy flow, and permanently inclusive of what computation cannot supply. These are not aspirational properties. They are necessary conditions, each load-bearing, each derivable from established theory.

The practical stakes are compressed by three dynamics. Path dependency makes the sequence of investment consequential and progressively less reversible. The visibility asymmetry between capability returns and fertility returns biases organizations toward the wrong investment first. And capability commoditization eliminates every form of AI advantage except the one that cannot be purchased: organizational fertility.

The architecture that satisfies these requirements is specified in Paper B. The failure mechanism it addresses is diagnosed in Paper A. This paper provides the theoretical foundations and generates six testable propositions - from the Fertility Index to the autonomy-to-initiative transition - that would move this framework from theoretical argument to empirical programme.

The medium is already there. Making it explicit, persistent, and available to computational agents is the central engineering challenge of enterprise AI.

References

- Argyris, C. and Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley.
- Angwin, J., Larson, J., Mattu, S., and Kirchner, L. (2016). Machine bias. *ProPublica*, 23 May.

- Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press.
- Bai, Y., Kadavath, S., Kundu, S., Askill, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., et al. (2022). Constitutional AI: harmless from AI feedback. arXiv:2212.08073.
- Christiano, P. F., Leike, J., Brown, T., Marber, M., Legg, S., and Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in Neural Information Processing Systems*, 30.
- Collins, H. (2010). *Tacit and Explicit Knowledge*. University of Chicago Press.
- Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1), pp. 128–152.
- Davenport, T. H. and Prusak, L. (1998). *Working Knowledge: How Organizations Manage What They Know*. Harvard Business School Press.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), pp. 332–337.
- Dalkir, K. (2011). *Knowledge Management in Theory and Practice* (2nd ed.). MIT Press.
- Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*, 10 October.
- Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), pp. 350–383.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), pp. 127–138.
- Garcez, A. d’A. and Lamb, L. C. (2023). Neurosymbolic AI: the 3rd wave. *Artificial Intelligence Review*, 56, pp. 12387–12406.
- Gartner (2025). Gartner predicts 40% of agentic AI projects will be scrapped by 2027. *Gartner Press Release*.
- Gartner (2025b). Lack of AI-ready data puts AI projects at risk. *Gartner Press Release*, 26 February 2025.
- Hollan, J., Hutchins, E., and Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), pp. 174–196.
- Hutchins, E. (1995a). *Cognition in the Wild*. MIT Press.
- Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19(3), pp. 265–288.
- Kauffman, S. A. (2019). *A World Beyond Physics: The Emergence and Evolution of Life*. Oxford Uni-

versity Press.

Kriegman, S., Blackiston, D., Levin, M., and Bongard, J. (2020). A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 117(4), pp. 1853–1859. <https://doi.org/10.1073/pnas.1910837117>

Kirchhoff, M., Parr, T., Palacios, E., Friston, K., and Kiverstein, J. (2018). The Markov blankets of life: autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15(138).

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), pp. 71–87.

Levin, M. (2019). The computational boundary of a “self”: developmental bioelectricity drives multicellularity and scale-free cognition. *Frontiers in Psychology*, 10, 2688.

Levin, M. (2022). Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 768201.

Luhmann, N. (1995). *Social Systems*. Stanford University Press.

Mingers, J. (1995). *Self-Producing Systems: Implications and Applications of Autopoiesis*. Plenum Press.

Maturana, H. R. (2002). Autopoiesis, structural coupling and cognition: a history of these and other notions in the biology of cognition. *Cybernetics & Human Knowing*, 9(3–4), pp. 5–34.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel Publishing.

McGilchrist, I. (2009). *The Master and His Emissary: The Divided Brain and the Making of the Western World*. Yale University Press.

McGilchrist, I. (2021). *The Matter With Things: Our Brains, Our Delusions, and the Unmaking of the World*. Perspectiva Press.

McGilchrist, I. (2024). Resist the machine apocalypse. *First Things*, March 2024.

McKinsey (2025). *The state of AI in 2025*. McKinsey Global Institute.

MITRE (2024). *AI Maturity Model*. The MITRE Corporation. <https://aidas.mitre.org/>

Nonaka, I. and Takeuchi, H. (1995). *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press.

Nonaka, I., Toyama, R., and Konno, N. (2000). SECI, Ba and Leadership: a unified model of dynamic knowledge creation. *Long Range Planning*, 33(1), pp. 5–34.

Ramstead, M. J. D., Badcock, P. B., and Friston, K. J. (2018). Answering Schrödinger’s question: a free-energy formulation. *Physics of Life Reviews*, 24, pp. 1–16. <https://doi.org/10.1016/j.plrev.2017.09.001>

- Precisely and Drexel University (2025). *The state of data quality for AI 2025*. Precisely, Inc.
- Polanyi, M. (1966). *The Tacit Dimension*. University of Chicago Press.
- Prigogine, I. (1980). *From Being to Becoming: Time and Complexity in the Physical Sciences*. W. H. Freeman.
- Prigogine, I. and Stengers, I. (1984). *Order Out of Chaos: Man's New Dialogue with Nature*. Bantam Books.
- Reichhart, W. and Gelas, A. (2026a). Dynamics blindness: when AI is locally correct and globally non-compliant. *SSRN Working Paper*.
- Gelas, A. and Reichhart, W. (2026b). The predictive organization: architecture for enterprise intelligence. *SSRN Working Paper*.
- Gelas, A. and Reichhart, W. (2026d). Governed intelligence architecture for institutional AI. *SSRN Working Paper*.
- Reichhart, W. and Gelas, A. (2026e). From autonomy to initiative: enterprise AI's real endgame. *SSRN Working Paper*.
- Snowden, D. J. and Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), pp. 68–76.
- Stacey, R. D. (2001). *Complex Responsive Processes in Organizations: Learning and Knowledge Creation*. Routledge.
- S&P Global (2025). 42% of companies scrapped most AI initiatives in 2025. *S&P Global Market Intelligence*.
- Varela, F. J. (1979). *Principles of Biological Autonomy*. North Holland.
- Veissière, S. P. L., Constant, A., Ramstead, M. J. D., Friston, K. J., and Kirmayer, L. J. (2020). Thinking through other minds: a variational approach to cognition and culture. *Behavioral and Brain Sciences*, 43, e90. <https://doi.org/10.1017/S0140525X19001213>
- Zhang, J. and Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive Science*, 18(1), pp. 87–122.
- Sydow, J., Schreyögg, G., and Koch, J. (2009). Organizational path dependence: opening the black box. *Academy of Management Review*, 34(4), pp. 689–709.