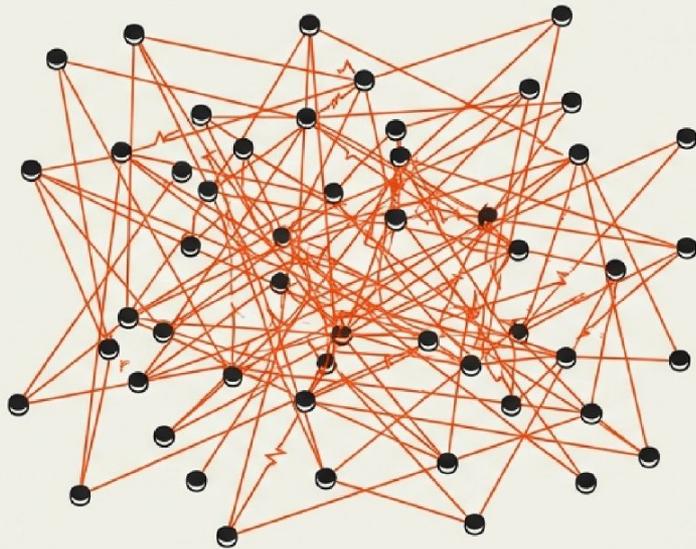
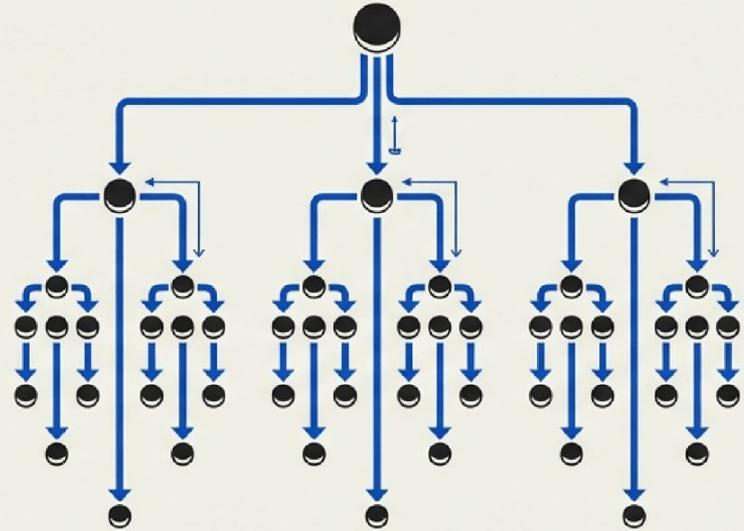


Organizational Theory as the Foundation for Agentic AI

MESH // CHAOTIC COORDINATION CRISIS



HIERARCHY // ENGINEERED STRUCTURE & SCALABILITY



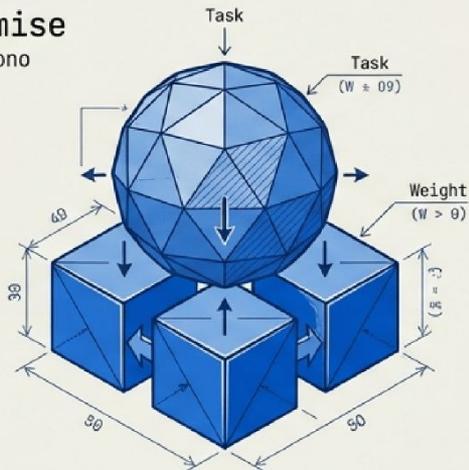
Why Multi-Agent Systems Fail at Scale and How Design Principles Solve the Coordination Crisis.

Based on "Organizational Theory as the Foundation for Agentic AI Systems" by Jonathan H. Westover, PhD.

The Coordination Tax: Why More Agents Doesn't Mean Better Results

The Promise

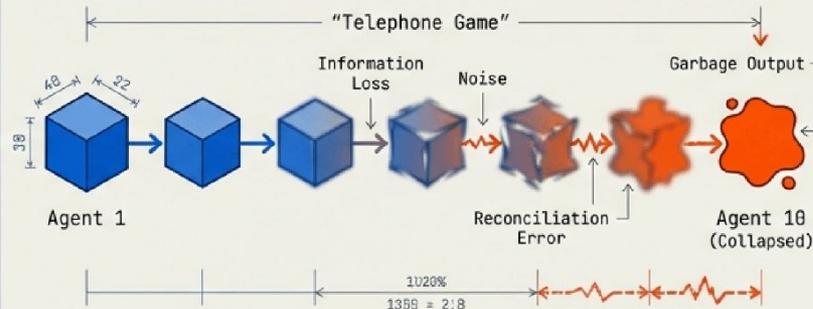
JetBrains Mono



Agentic AI is defined not as chatbots, but as systems that maintain state, plan sequences, and invoke tools. The expectation is “Agent Swarms” that self-organize to solve complex problems.

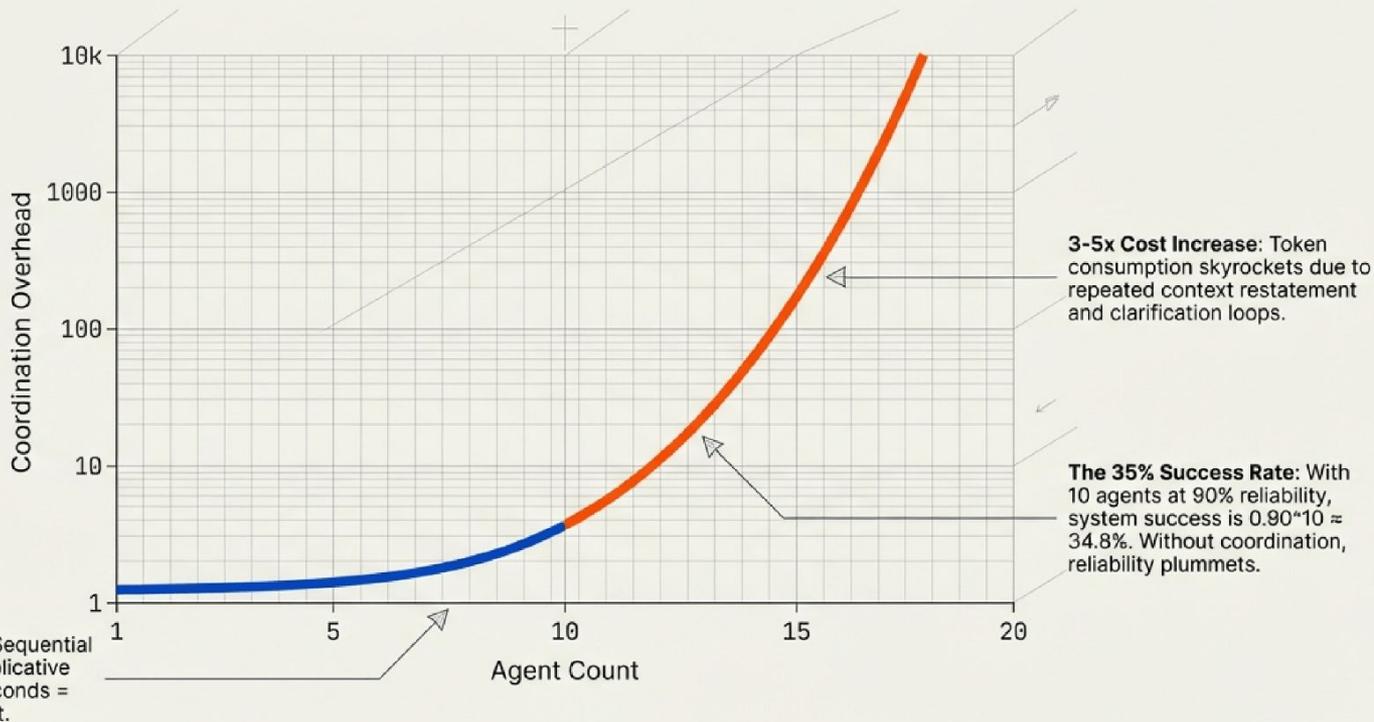
The Reality

JetBrains Mono



Systems that work elegantly with 3 agents collapse under overhead with 10. Information degrades as it passes between agents. Token costs explode as agents clarify, restate, and reconcile, creating a “Coordination Tax” that kills ROI.

The Mathematics of Failure



The bottleneck isn't model capability. It is coordination.

We Are Repeating the Management Mistakes of 1920

Organizational theorists (Galbraith, March & Simon) have studied coordination for a century. We are currently building AI systems that assume models have 'unlimited managerial capacity.' They do not.

	Human Org Concept (1920s)	AI Architecture Reality (2020s)
Span of Control (Managerial Attention)	Departmental Silos	Context Window & Attention Limits
Memos & Specifications		Specialized Agent Roles
Bounded Rationality (Cognitive Limits)		Structured Artifacts (JSON/Schemas)
		"Lost in the Middle" Phenomenon

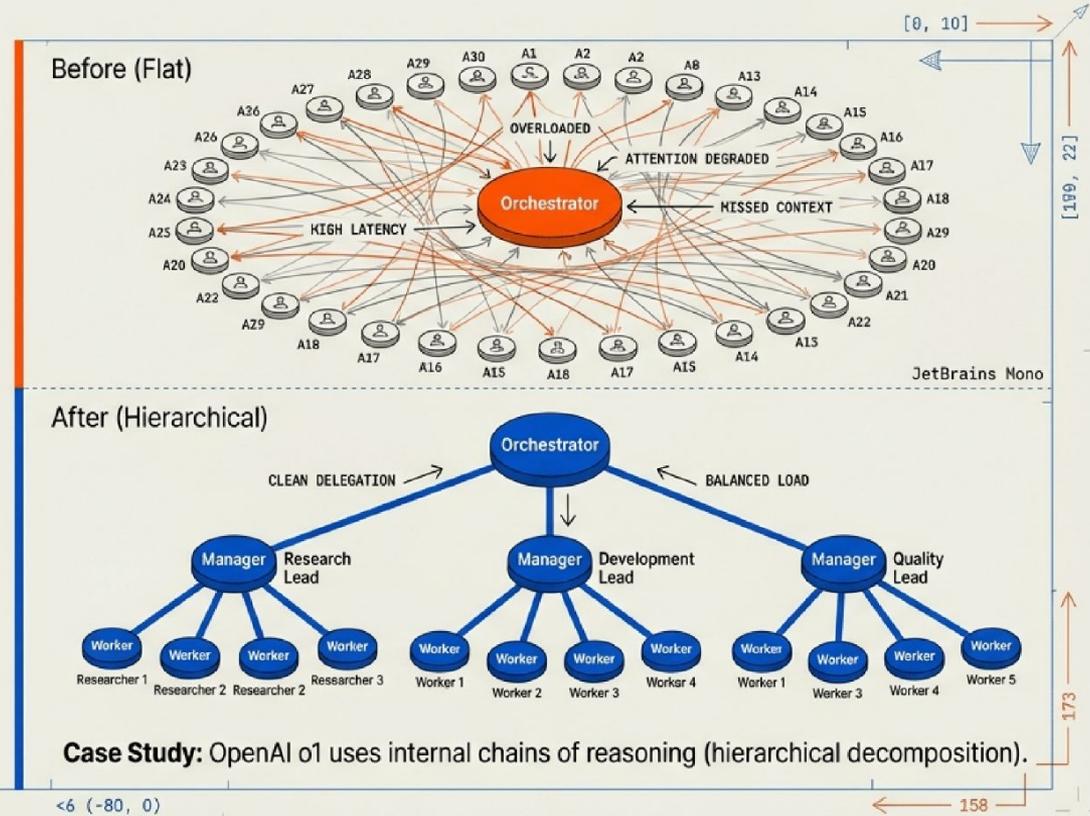
“Organizations that recognize agentic AI as fundamentally an organizational design problem stand to gain significant competitive advantage.”

Principle 1: Hierarchy and Span of Control

The Problem: An 'Orchestrator' agent cannot effectively manage 50 subordinates. Attention degrades.

The Fix: Classical theory (Urwick, 1956) suggests a span of control of 5–7 direct reports.

AI Design Pattern: Introduce 'Middle Manager' agents. Group agents into functional clusters (e.g., Research Lead coordinates 5 researchers).

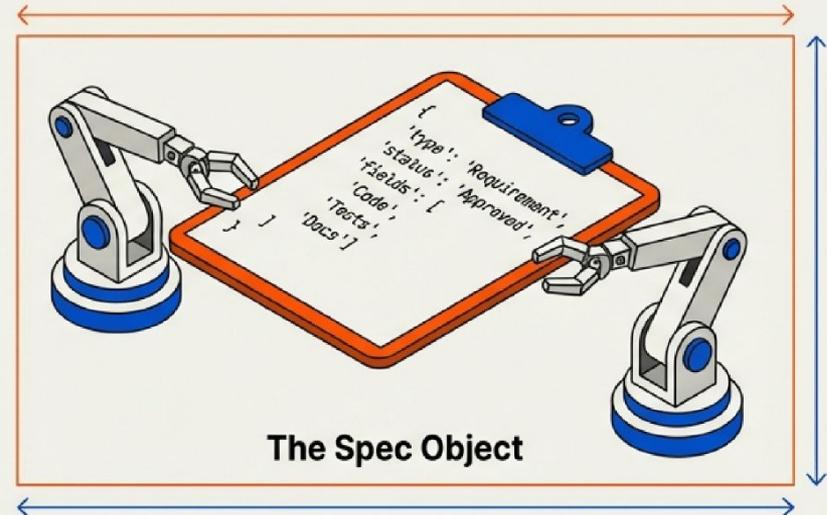
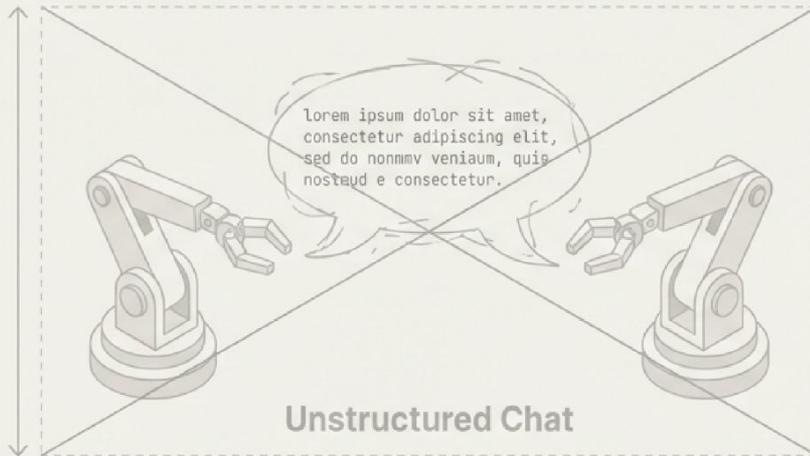


Principle 2: Boundary Objects Over Unstructured Chat

The Problem: Passing freeform text between agents creates the “telephone game” and hallucination.

The Fix: Boundary Objects (Star & Griesemer, 1989). Shared artifacts that different groups can reference.

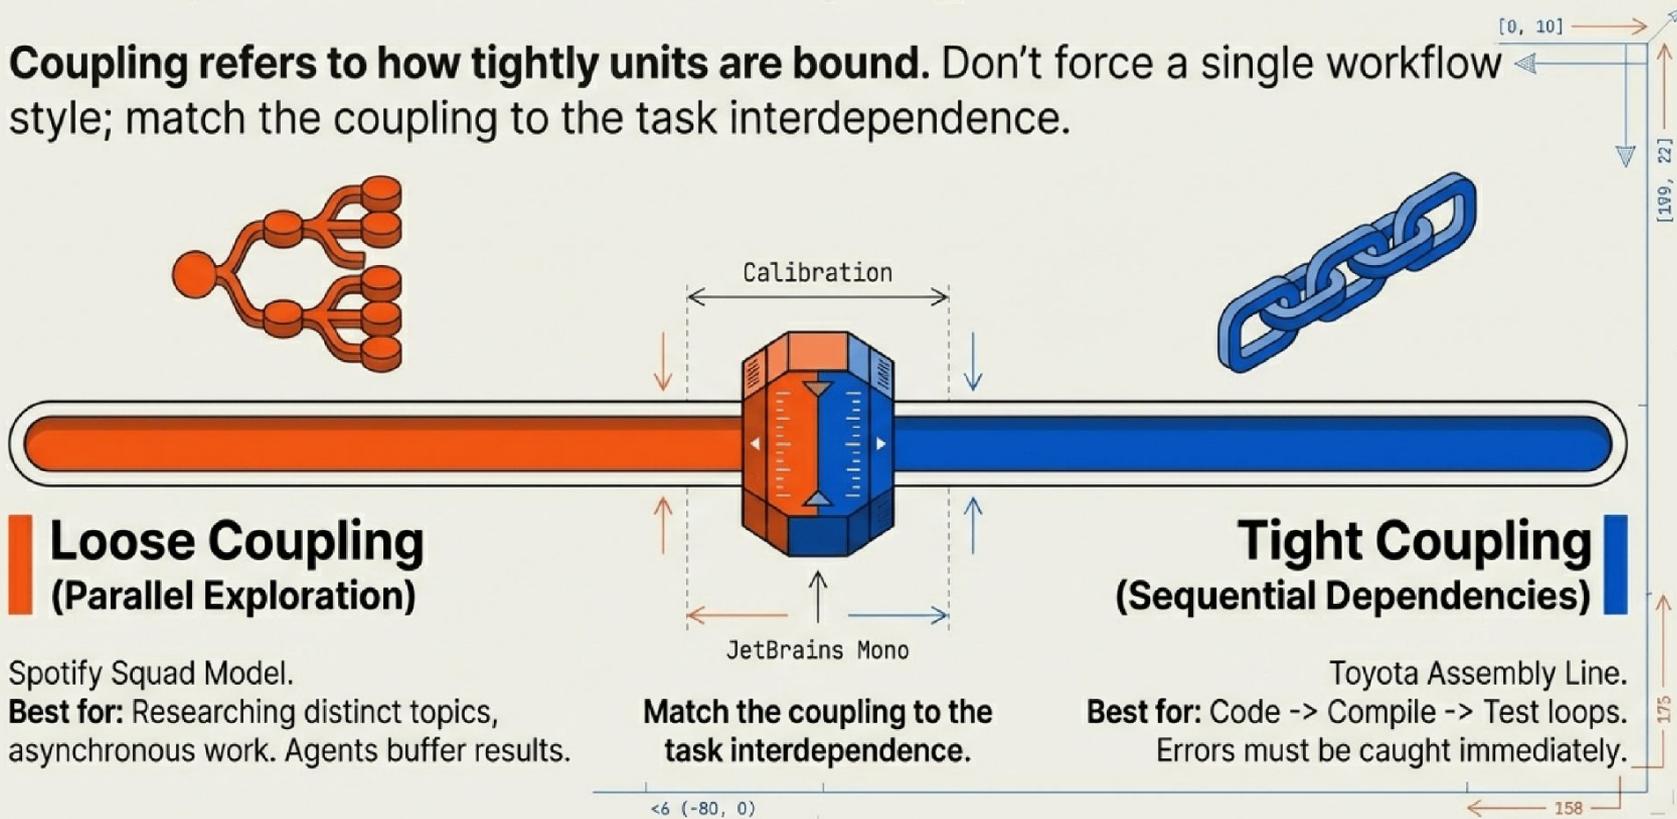
JetBrains Mono



Industry Parallel: Just as Stripe uses detailed API specs to align teams, AI agents need structured schemas (Requirements Object, Design Spec Object) to align research and execution.

Principle 3: Calibrated Coupling

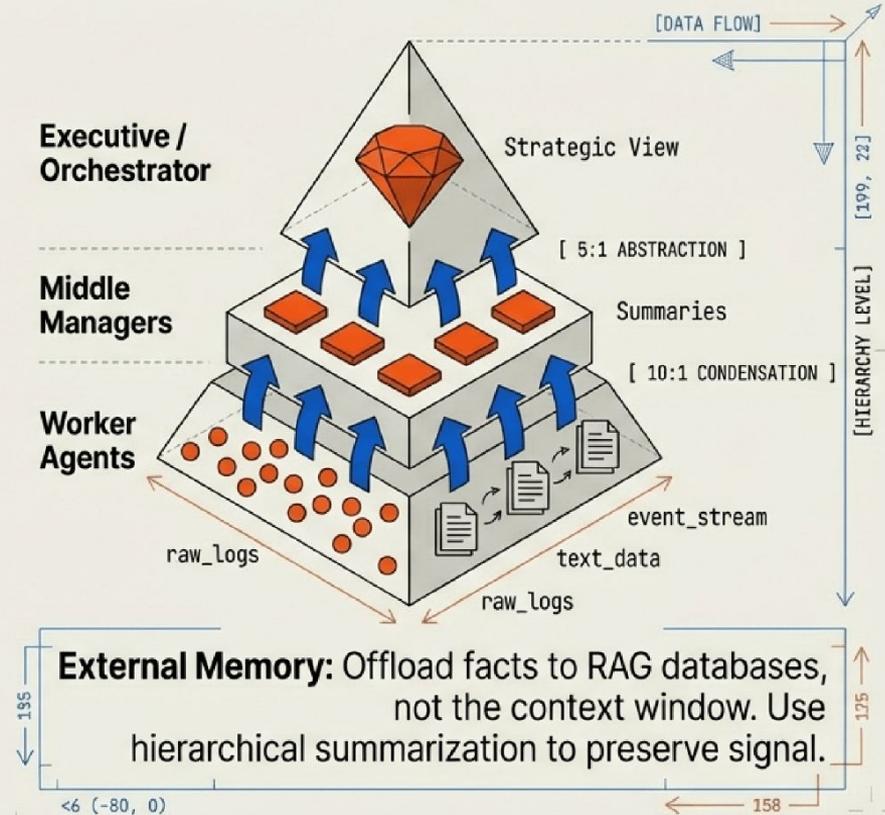
Coupling refers to how tightly units are bound. Don't force a single workflow style; match the coupling to the task interdependence.



Principle 4: Managing Bounded Rationality

The Problem: “Lost in the Middle” phenomenon. Models have finite attention. Feeding an agent the entire project history ensures failure.

The Fix: Division of labor and information abstraction (Simon, 1957). Managing an agent serups in the uli- primate project abstraction, chariflot toforth middle managers. Fine point mustcenate the agents loss analysis are not matamap.



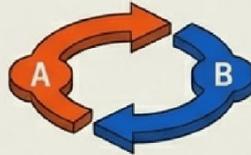
Principle 5: Shared Language and Protocols

Even with specs, interpretive disconnects occur. “Improve code” means different things to different models. We need strict protocols.



Domain Ontologies

Define strict vocabularies. Example: ‘Critical’ vs ‘High Priority’ must be mathematically defined.



Read-Back Protocols

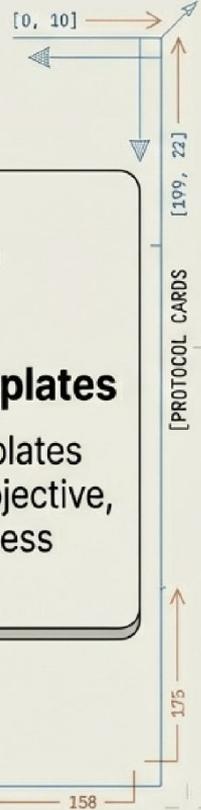
Agent B must restate the understanding of the task to Agent A before executing. (Borrowed from Aviation).



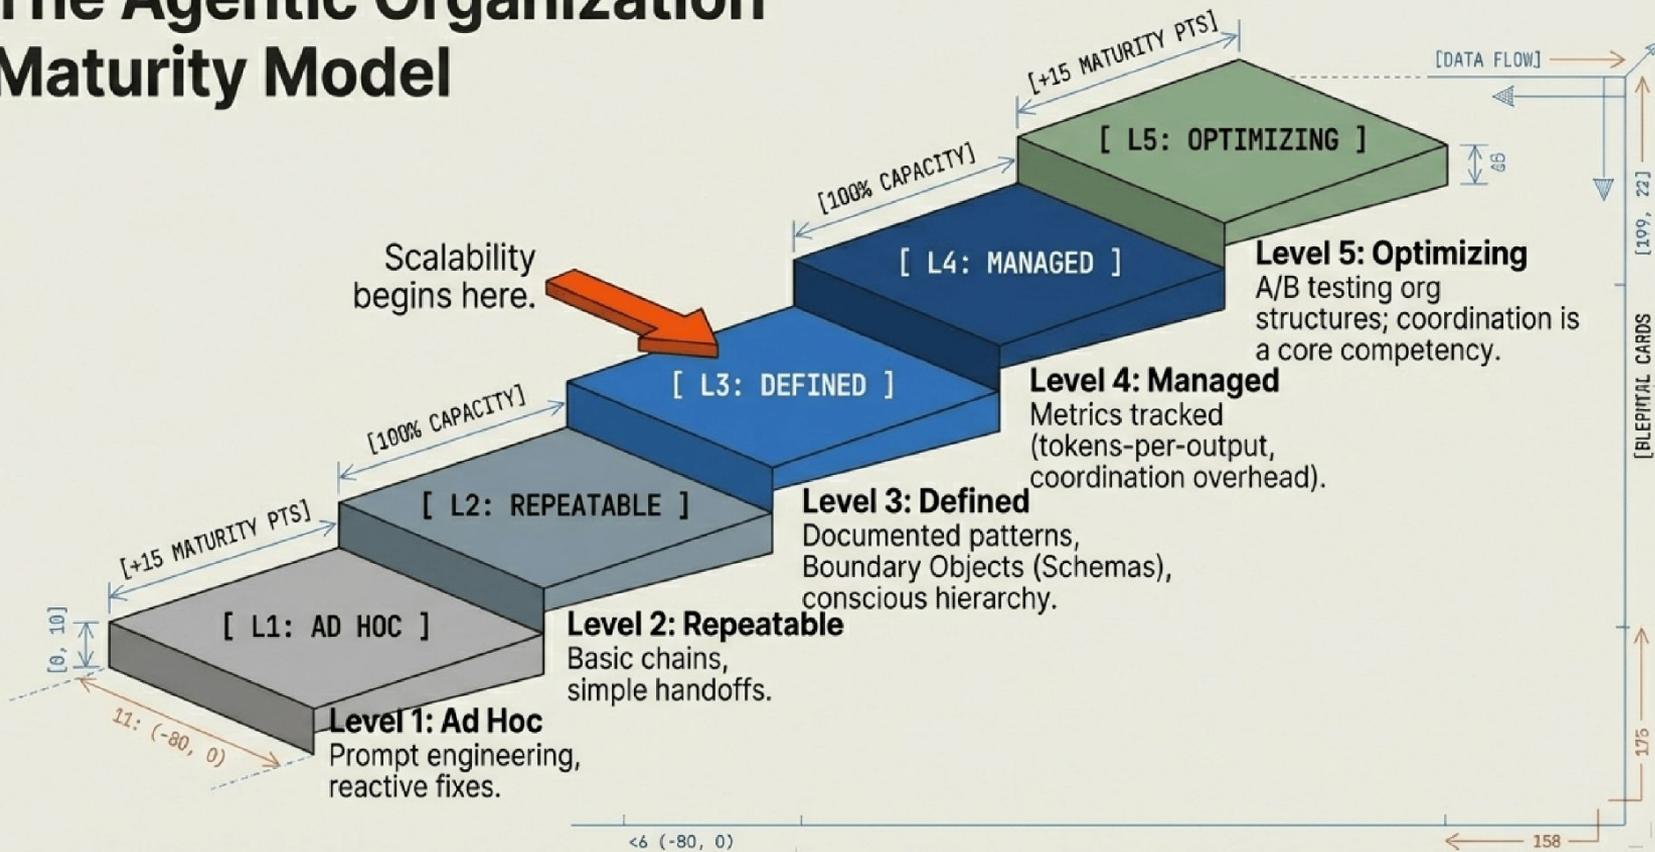
Standardized Templates

Task assignment templates with fixed fields for Objective, Constraints, and Success Criteria.

Industry Parallel: Boeing engineers use ‘interface control documents’ to prevent physical crashes. AI needs the same to prevent logic crashes.



The Agentic Organization Maturity Model



Building the Cross-Disciplinary Design Team

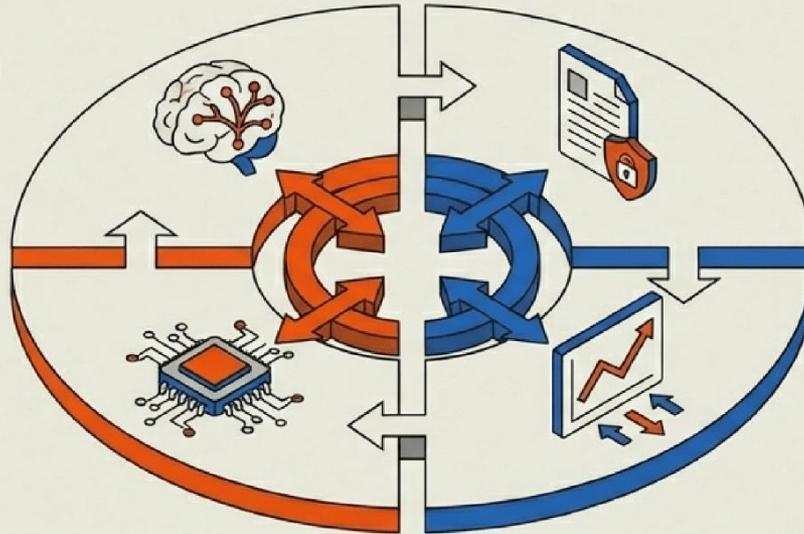
You cannot build Level 3+ systems with ML Engineers alone.

Org Psychologists

Design roles, flows, span of control.

AI Engineers

Implementation & model constraints.



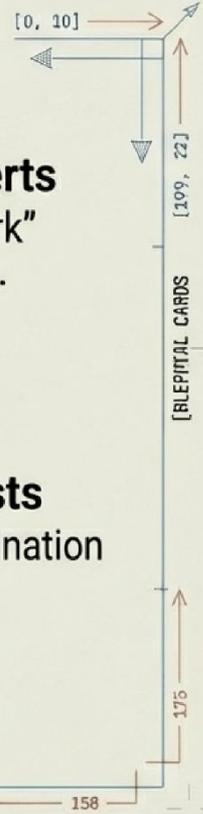
Domain Experts

Define the “work” and ontologies.

Data Scientists

Measure coordination efficiency.

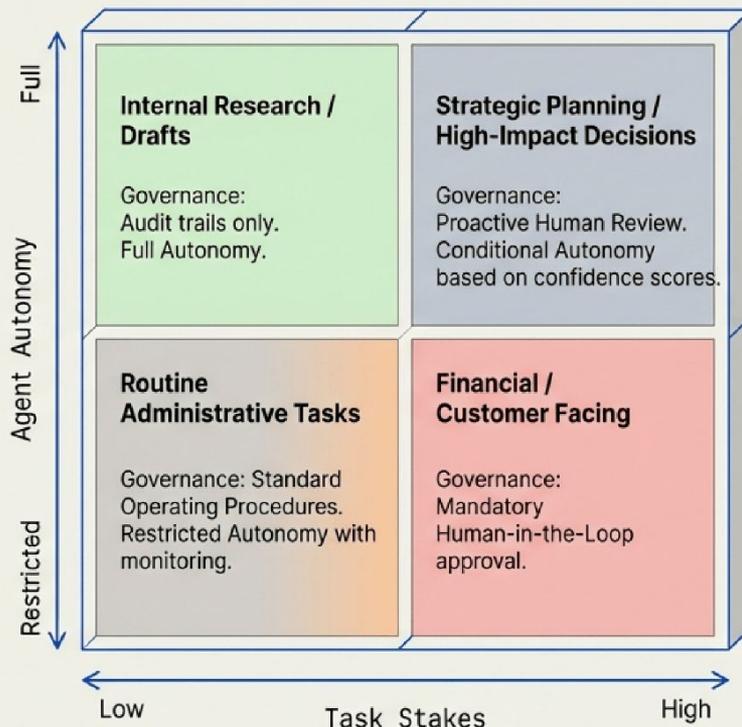
Treat the team as a ‘Connect + Develop’ ecosystem (P&G), integrating diverse expertise rather than siloing tech.



Governance: Balancing Autonomy with Safety

The governance framework must adapt to the risk profile of the task and the agent's capabilities. This approach ensures operational safety while maximizing efficiency and innovation.

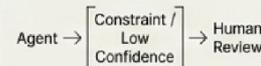
Risk-Based Oversight



Exception-Based Oversight Protocol

Define clear operating parameters.

If an agent hits a constraint or low confidence signal, it escalates to a human. This parallels corporate "Delegation of Authority".



[0, 10]

[JetBrains Mono]

[JetBrains Mono]

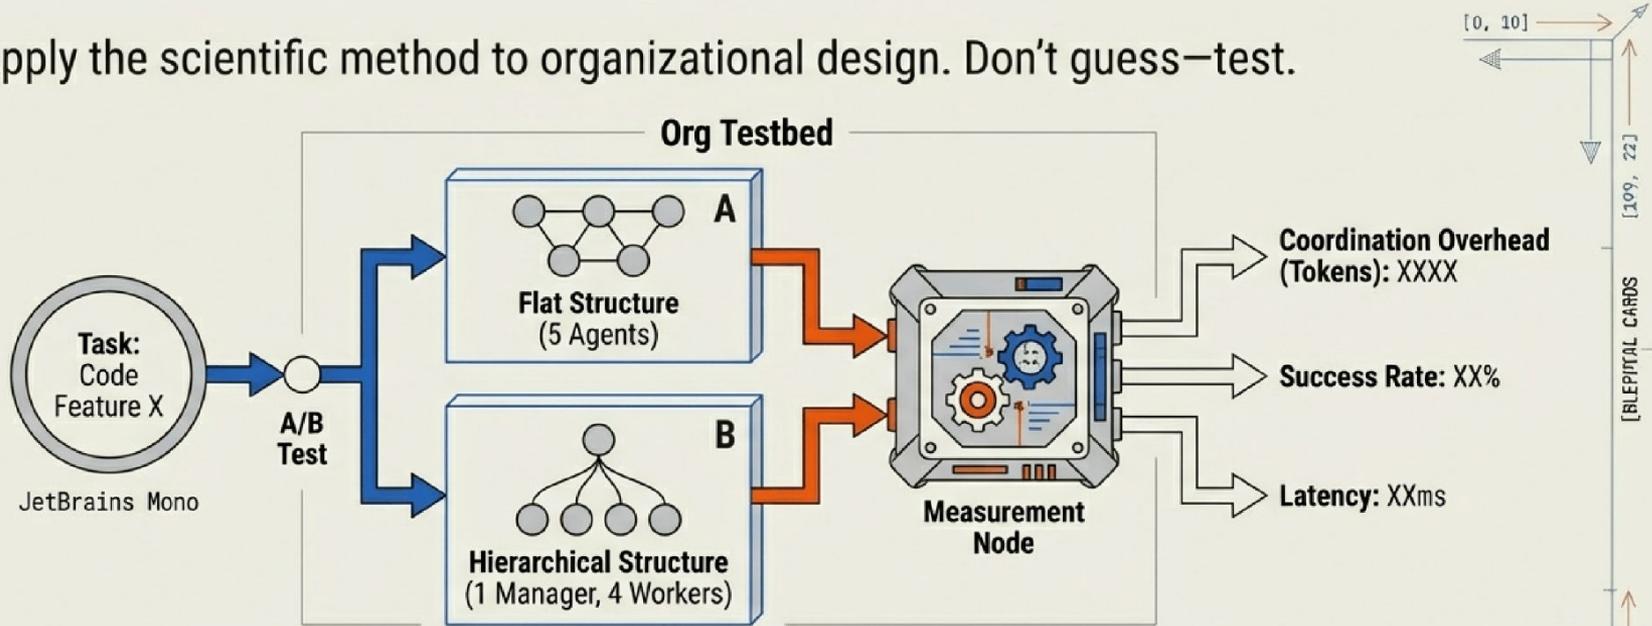
[19, 22]

[109, 22]

[BLEPTAL CARDS]

Infrastructure for Experimentation

Apply the scientific method to organizational design. Don't guess—test.



Action: Move from 'Prompt Engineering' to 'Organization Engineering'. Run simultaneous tests to identify optimal structures, referencing Google's experimentation culture.

Summary of Design Patterns

Problem	Org Theory Solution	Technical Implementation
Span of Control Overload	Hierarchy	Intermediate "Manager" Agents
Ambiguity in Handoffs	Boundary Objects	JSON Schemas & Specs
Bottlenecks vs. Drift	Calibrated Coupling	Phase-based Async/Sync flows
Context Limits	Division of Labor	RAG & Hierarchical Summarization
Misinterpretation	Shared Ontologies	Read-back Protocols

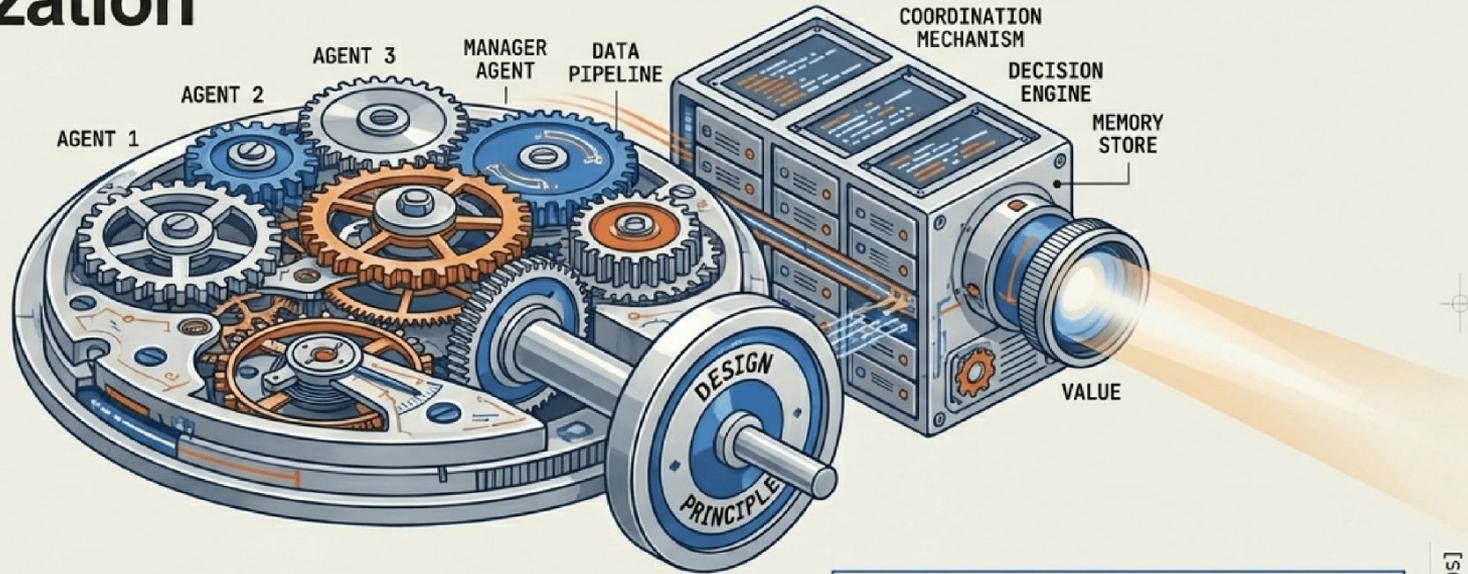
<2 (-250 , 1)>

[BLEPITAL CARDS]

<6 (-80 , 0)>

Domaine Text, Slate Grey

The Competitive Advantage is Organization Organization



Models are becoming commodities. The differentiating factor is not the intelligence of the individual model, but the organization of the system.

Stop treating agents like chat tools. Start treating them as an **organizational design challenge**.

<2 (-250 , 1)>

<6 (-80 , 0)>

[BLEPITAL CAROS]

[ORGANIZATIONAL CARD]