

Strategy in the Age of Agentic AI: Governing the Strategic Envelope

Executive Summary

Soumitra Dutta¹ and Yves Doz²

A customer contacts a global retailer about a missing order. In a traditional organization, the case moves through a familiar sequence: a service agent reviews the order, contacts logistics, perhaps escalates to a supervisor, and eventually proposes a resolution. In an agentic organization, something different happens. An AI service agent identifies the customer, checks shipment status, evaluates the customer's history and value, queries inventory, assesses alternative fulfilment options, and proposes a response — all in seconds. A logistics agent considers rerouting. A risk agent checks for fraud patterns. If the situation falls within defined boundaries, the system acts. If it crosses a threshold — high-value customer, legal exposure, emotional escalation — a human specialist is brought in.

No single manager directed this sequence. No central decision-maker saw the interaction. Yet the organization acted. This is the new reality created by agentic AI — systems that can reason, pursue goals over time, coordinate with other systems, and operate with limited human intervention. Individually, these systems are powerful. Collectively, they change the nature of the organization itself. And that change creates a challenge that most leaders have not yet confronted: how do you govern what your organization becomes when thousands of AI agents are acting, adapting, and interacting continuously?

From Tools to Populations of Agents

For most of the past decade, executives thought about AI as a tool. It helped analysts detect patterns, marketers personalize campaigns, and employees automate routine work. Agentic AI changes that framing entirely. These systems do not merely support decisions. They make them. They monitor conditions, initiate actions, coordinate across functions, and refine their behavior based on outcomes — and they can be replicated at very low cost across the organization.

A supply chain may contain agents that monitor suppliers, detect disruptions, adjust inventory, and recommend rerouting. A bank may deploy agents that screen transactions, monitor risk exposure, and escalate anomalies. A customer service organization may rely on agents that resolve routine cases, personalize interactions, and determine when human judgment is required. As these systems multiply, something fundamental changes. Decision-making becomes continuous rather than episodic, distributed rather than centralized, and increasingly system-driven rather than manager-driven.

¹ Director, Portulans Institute, Washington DC and Former Founding Dean, SC Johnson College of Business, Cornell University, New York, USA

² Solvay Chaired Professor of Technology Innovation Emeritus, INSEAD, France

Organizations as Complex Adaptive Systems

As agentic systems proliferate, organizations begin to resemble complex adaptive systems more than traditional hierarchies. A pricing decision affects demand. Demand affects inventory. Inventory affects fulfilment. Fulfilment affects customer satisfaction. Satisfaction affects future demand. Each local decision changes the environment in which other decisions are made.

This kind of emergence — where patterns of organizational behavior arise from the interaction of many actors — has always existed. What is new is the scale, speed, and autonomy of emergence in agentic organizations. Emergence is no longer occasional. It becomes routine. And when emergence becomes routine, strategy becomes a discipline of governing it.

Bounded Autonomy: The Engineering Insight

Engineers have long understood that complex systems cannot be controlled action by action. In aviation, robotics, and distributed computing, systems move too quickly and contain too many interdependencies for direct supervision to work. The solution is not to eliminate autonomy. It is to bound it. A modern aircraft is not manually controlled at every moment. Automated systems manage most of the flight, but within defined limits. The pilot supervises, interprets unusual conditions, and intervenes when the system approaches the boundaries of its operating envelope.

Agentic organizations require the same logic. Leaders cannot specify every action or predict every interaction among human and AI agents. But they can define the boundaries within which action is allowed to emerge. The managerial challenge is not to suppress autonomy. It is to make autonomy safe enough to scale.

The Strategic Envelope: Purpose, Constraints, and Feedback

This is the role of what the paper calls the strategic envelope: the range of actions, behaviors, and adaptations within which the organization's human and agentic systems are permitted to operate. A well-designed strategic envelope has three elements.

Purpose defines direction. Leaders must specify what the organization is trying to achieve, what trade-offs matter, and what values must be preserved. In a distributed agentic system, purpose is not merely aspirational — it is the organizing logic that allows thousands of local actions to contribute to a coherent whole.

Constraints define boundaries — what must not happen, what risks are unacceptable, and when systems must defer to human judgment. A bank may allow agents to recommend products but impose strict constraints on how vulnerable customers are treated. A retailer may allow dynamic pricing but set limits to prevent reputational damage. A healthcare system may support AI diagnosis but require human review before high-stakes treatment decisions.

Feedback maintains alignment. Even well-designed systems drift. Markets change, data shifts, and agents learn from new patterns in ways leaders did not anticipate. Performance metrics, anomaly detection, escalation rules, and audit processes allow the organization to detect when behavior is moving outside the envelope before drift becomes damage.

Together, purpose, constraints, and feedback create the conditions for bounded emergence. The system can adapt — but not in any direction. The envelope's purpose is not to slow the organization down. It is to make autonomy safe enough to scale.

Intended Strategy vs. Produced Strategy

In agentic organizations, the patterns that actually emerge may not reflect leadership intent. A company may intend to compete on trust, but its recommendation agents may gradually optimize toward short-term engagement. It may intend to build supply-chain resilience, but its agents may repeatedly favor cost and speed. It may intend to improve customer relationships, but its service agents may prioritize case-closure rates over relationship quality.

This gap — between intended strategy and produced strategy — is where strategic drift begins. The produced strategy is not what leaders announced. It is the pattern of behavior generated by the interaction of human and AI agents over time. Reading that pattern, interpreting what it reveals about the direction the organization is actually moving, and adjusting the envelope when needed: this is one of the most important new strategic tasks of the agentic era.

Leadership as Envelope Governance

Agentic AI changes where leadership is exercised. Leaders who still try to direct action from the top will find themselves overwhelmed and increasingly irrelevant. Instead, leaders become governors of the strategic envelope — the conditions under which autonomous action emerges.

This role has four elements. Leaders define purpose: what is worth optimizing for is a human judgment that no system can make. They calibrate autonomy: not every part of the organization should have the same degree of freedom — routine decisions may be fully automated, while ethical and regulatory questions require human review. They interpret emergence: reading the patterns that reveal whether the organization is moving in the intended direction or drifting away from it. And they intervene selectively: not through constant override, but when boundaries are breached, new patterns emerge, or the envelope itself needs redesigning.

This resembles what military strategists call mission command — leaders define intent and limits while local actors adapt to changing conditions. The difference is that in agentic organizations, many of those local actors are not people. They are intelligent systems operating at a scale and speed that no hierarchy can match.

What Remains Human

The more agentic the organization becomes, the more important human judgment becomes in certain domains. Systems can optimize, adapt, and recognize patterns. But they cannot determine purpose, legitimacy, or meaning. Leaders must decide which trade-offs are acceptable, interpret ambiguous signals, and remain accountable for outcomes — even when no individual human made the specific decision that produced them.

Strategy in the age of agentic AI is not a technical problem. It is a deeper leadership problem. Small choices about objectives, constraints, and feedback can shape millions of actions. The organization will adapt. The question is whether it adapts coherently — toward the outcomes leadership intends — or drifts toward something else entirely.

The organizations that succeed in this environment will not be those that eliminate autonomy or suppress emergence. They will be those that govern it — so that even as behavior evolves continuously, it remains aligned, responsible, and strategically coherent.

Disclosure: AI was used research and copy editing of this article.