

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

S. Dutta and Y. Doz

*Former Founding Dean, SC Johnson College of Business, Cornell University, USA
Solvay Chaired Professor of Technology Innovation Emeritus, INSEAD, France*

A customer contacts a global retailer because an order has not arrived. In a traditional organization, the request would move through a familiar sequence: customer service reviews the case, checks inventory, contacts logistics, perhaps escalates to a supervisor, and eventually proposes a resolution.

In an agentic organization, the interaction unfolds differently. An AI service agent identifies the customer, reviews the order history, checks shipment status, evaluates the customer's value and prior experience, queries inventory availability, identifies alternative fulfillment options, and proposes a resolution. A logistics agent assesses whether replacement inventory can be rerouted. A pricing or policy agent determines whether compensation is appropriate. A risk agent evaluates whether the request fits known fraud patterns. If the situation remains within defined boundaries, the system acts. If it crosses a threshold—high-value customer, legal exposure, repeated failure, emotional escalation—it routes the case to a human specialist.

No single manager directs this sequence. No central decision-maker sees every interaction. Yet the organization acts. This is the new strategic reality created by agentic AI. Artificial intelligence is moving beyond analysis, recommendation, and task automation. Increasingly, AI systems can think, reason, pursue goals over time, sequence actions, coordinate with other systems, adapt to feedback, and operate with limited human intervention.

Individually, these systems are powerful. Collectively, they change the nature of the organization. The firm begins to resemble a complex adaptive system: a network of human and artificial agents interacting continuously, making local decisions, and producing patterns of behavior that no single actor fully controls.

The challenge for strategy is no longer only to decide what the organization should do. It is to govern what the organization becomes as thousands or more of human and agentic decisions interact over time. Programmable strategy explains how intent is translated into system behavior. The strategic envelope addresses the next problem: how to govern what happens when many human and AI agents act, adapt, and interact within that system.

From Tools to Populations of Agents

For much of the last decade, executives thought about AI as a tool. It helped analysts detect patterns, marketers personalize campaigns, managers forecast demand, and employees automate routine work.

Agentic AI changes that framing. These systems do not simply support decisions. They contain relevant knowledge, can reason and decide how to participate in organizational activity. They can monitor conditions, initiate action, coordinate across systems, and refine their behavior based on outcomes. More importantly, they can proliferate. Once developed, agents can be replicated across functions, geographies, and workflows at very low marginal cost.

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, adjust credit exposure, and escalate anomalies. A customer service organization may rely on agents that resolve routine issues, personalize responses, and identify when human intervention is required.

As these systems multiply, the organization changes. It becomes a distributed system of interacting agents—some human, many agentic. Decision-making becomes continuous rather than episodic. It becomes distributed rather than centralized. And it becomes increasingly system-driven rather than manager-driven.

This does not mean humans disappear. But it does mean that humans increasingly intervene selectively: when ambiguity is high, when constraints are breached, when values are at stake, or when the system encounters situations outside its designed envelope.

Organizations as Complex Adaptive Systems

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

The same logic applies in risk management, customer service, marketing, logistics, and platform governance. The behavior of the organization emerges from interactions across hundreds and thousands of interconnected actions and decisions.

This is not entirely new. Strategy scholars have long recognized that realized strategy often differs from intended strategy. Henry Mintzberg described strategy as emerging from patterns of action over time. Sumantra Ghoshal emphasized the importance of guided emergence in complex organizations. Bower and Burgelman showed that strategy is shaped not only by top-down intent, but also by resource allocation processes, middle-level integration, and bottom-up variation.

What is new is the scale, speed, and autonomy of emergence. In earlier organizations, emergence came largely from human initiative, local adaptation, and managerial interpretation. In agentic organizations, emergence is increasingly produced through interactions across a multitude of AI agents (and some humans) operating continuously, in real time and at scale. The implication is profound: emergence is no longer occasional. It becomes routine. And when emergence becomes routine, strategy becomes a discipline of governing it.

What Engineering Teaches Us About Bounded Autonomy

Engineers have long understood that complex systems cannot be controlled action by action. In aviation, robotics, power grids, and distributed computing, systems move too quickly, contain too many interdependencies, and face too many possible conditions 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 much of the flight, but they operate within defined limits. The pilot does not direct every adjustment; instead, the pilot supervises, interprets unusual conditions, and intervenes when the system approaches situations that are ambiguous, unsafe, or outside its operating envelope.

The same principle appears in power grids and distributed computing systems. Local components adjust continuously, but they do so within constraints that prevent cascading failure. The goal is not perfect prediction of every possible event. It is designing a system that can adapt without becoming unstable.

Agentic organizations require the same logic. Leaders cannot specify every action an AI-enabled organization should take, nor can they predict every interaction among human and artificial agents. But they can define the boundaries within which action is allowed to emerge.

This is the principle of bounded autonomy. Systems are given room to act, adapt, and coordinate, but only within limits that preserve safety, trust, and strategic coherence. The managerial challenge is not to suppress autonomy. It is to make autonomy safe enough to scale.

Strategy as a Strategic Envelope

In traditional strategy, the central question was often: What should we do? In an agentic organization, that question remains important, but it is no longer sufficient. A more important question emerges: How do we ensure that what the organization does—through the interaction of human and agentic systems—remains aligned with what we intend?

This is the role of the strategic envelope. The strategic envelope defines the range of actions, behaviors, and adaptations within which the organization's human and agentic systems are allowed to operate. It does not prescribe every decision. It defines the space within which decisions can emerge.

A well-designed strategic envelope has three elements: purpose, constraints, and feedback.

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 becomes the organizing logic that allows many local actions to contribute to a coherent whole.

Constraints define boundaries. They specify 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 avoid reputational damage or perceived unfairness. A healthcare system may use AI to support diagnosis, but require human review before high-stakes treatment decisions.

Feedback maintains alignment. Even well-designed systems drift. Markets change, data shifts, customers respond in unexpected ways, and agents learn from new patterns. Feedback mechanisms—performance metrics, anomaly detection, escalation rules, human review, and audit processes—allow the organization to detect when behavior is moving outside the envelope.

Together, purpose, constraints, and feedback create the conditions for bounded emergence. The system can adapt. But it cannot adapt in any direction. The envelope must be enabling, not merely restrictive. Its purpose is not to slow the organization down, but to make autonomy safe enough to scale.

Strategy Becomes What the System Produces

Within the strategic envelope, behavior is not centrally dictated. It emerges. Consider a retailer using agentic systems across pricing, inventory, promotion, and customer engagement. Each system optimizes locally. Pricing agents respond to demand and competitor moves. Inventory agents reallocate stock. Promotion agents test offers. Customer service agents adapt interactions. No single decision defines the company's strategy. But over time, patterns emerge. The firm may become more price-aggressive, more service-oriented, more resilient, or more volatile. These patterns may reflect leadership intent — or they may reveal drift.

In agentic organizations, leaders must therefore compare intended strategy with produced strategy. Intended strategy is what leaders define: the priorities, trade-offs, and values they want the organization to pursue. Produced strategy is the pattern of behavior generated by the interaction of human and AI agents over time. The gap between the two is where strategic drift potentially begins.

This is why agentic AI changes the work of strategy. A company may intend to compete on trust, but its recommendation agents may gradually optimize toward engagement. It may intend to build resilience, but its supply-chain agents may repeatedly favor cost and speed. It may intend to improve customer experience, but its service agents may prioritize closure rates over relationship quality.

The same logic applies across functions. In supply chains, repeated agentic adjustments to sourcing, logistics, and production can gradually shift the firm from a cost-minimization posture to a resilience posture — or the reverse. In customer engagement, millions of AI-mediated interactions can shape brand perception more powerfully than a formal marketing campaign. In risk management, small local adjustments can accumulate into a more conservative or more aggressive organizational posture.

In an agentic organization, strategy is therefore not only what leaders intend. It is what the system consistently produces. The strategic task is to read those patterns, interpret what they reveal, and adjust the envelope when necessary.

Leaders must ask: What patterns are forming? Are they aligned with our purpose? Where are local optimizations creating system-level drift? Are we becoming the organization we intended to become?

Strategy, in this context, becomes less about checking whether individual decisions complied with a plan and more about understanding what the system is becoming through its own behavior. That is one of the most important new leadership tasks in the age of agentic AI.

Leadership as Envelope Governance

Agentic AI does not eliminate leadership. It changes where leadership is exercised. Leaders can no longer rely primarily on directing action from the top. The scale and speed of agentic decision-making make that impossible. Instead, leaders become governors of the strategic envelope. Leaders become governors of the conditions under which autonomous action emerges.

Their role has four parts. First, they define purpose. This is the human foundation of strategy. Systems can optimize, but they cannot decide what is worth optimizing for.

Second, they calibrate autonomy. Not every part of the organization should have the same degree of freedom. Routine customer inquiries may be handled autonomously. Ethical issues, regulatory questions, and high-stakes exceptions may require human judgment.

Third, leaders interpret emergence. They must learn to read the system—to identify patterns that reveal whether the organization is moving in the intended direction or drifting away from it.

Fourth, they intervene selectively. In agentic organizations, intervention should not mean constant override. It should occur when boundaries are breached, when new patterns emerge, or when the envelope itself needs to be redesigned.

This resembles what military strategists have long called mission command: leaders define intent and boundaries, 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.

Leadership becomes less about issuing instructions and more about shaping the conditions under which action occurs. Leaders are no longer only decision-makers or system designers. They are governors of the conditions under which autonomous action emerges.

When the Envelope Fails

Poorly designed strategic envelopes create predictable risks. If the envelope is too tight, the organization becomes brittle. Agentic AI and employees lack the freedom to respond to changing conditions. Innovation slows. Local adaptation is suppressed.

If the envelope is too loose, the organization becomes unstable. Systems optimize aggressively but without adequate boundaries. Pricing agents may damage trust. Engagement agents may maximize short-term interaction at the expense of long-term brand health. Risk agents may become too permissive or too conservative.

If the envelope is misaligned, the organization becomes incoherent. Different systems optimize for different goals. Some pursue growth, others minimize risk, others maximize efficiency. Each may perform well locally, while the organization moves in conflicting directions.

The danger is that the system may appear to work. It may act quickly, efficiently, and continuously. But it may still produce the wrong outcomes. In agentic organizations, the critical diagnostic question is no longer simply: Who made the wrong decision? It is: What system produced this behavior?

What Remains Human

The more agentic the organization becomes, the more important human judgment becomes. Systems can optimize. They can adapt. They can identify patterns. But they cannot determine purpose, legitimacy, or meaning.

Leaders must define what the organization stands for. They must decide which trade-offs are acceptable. They must interpret ambiguous signals. And they remain accountable for outcomes, even when no individual human made the specific decision. This is why agentic AI does not turn strategy into a technical problem. It turns strategy into a deeper leadership problem.

When behavior is generated continuously across distributed systems, the consequences of design choices are amplified. Small choices about objectives, constraints, or feedback can shape thousands or millions of actions. The organization will adapt. The question is whether it will adapt coherently.

Governing Bounded Emergence

The rise of agentic AI does not eliminate strategy. It transforms it. Strategy is no longer only a plan, a position, or a set of choices. It is the governance of behavior in a system that learns, adapts, and acts continuously.

The central challenge is no longer to specify every move. It is to define the space within which moves can be made. That space is the strategic envelope.

In this world, leaders must think less like commanders issuing instructions and more like designers of complex systems. They must define purpose clearly, establish enabling constraints, build feedback mechanisms, interpret emerging patterns, and intervene when the system drifts.

Strategy becomes the art of governing bounded emergence. The organizations that succeed will not be those that eliminate autonomy or suppress emergence. They will be those that guide it—so that, even as behavior evolves, it remains aligned, responsible, and strategically coherent.

Disclosure: AI was used for research and copy-editing while writing this paper.