As AI Gets Better, Strategy Gets Harder

Max Michaels | Michael Leiblein | Rao Mikkilineni

There is a peculiar grandeur in the way companies now talk about AI. They speak of acceleration, scale, augmentation, as though the future will belong simply to those who can move first and move most. Yet the more revealing drama is quieter. AI has not merely made organizations better at seeing. It has made them capable of acting on far more than they can fully understand. That is the real shift. The old strategic bottleneck was the generation of options. The new one is the governance of commitment. Once variation begins to outrun the capacity to verify and judge, the firm enters a different landscape altogether, one in which the danger is not blindness but overproduction of action, not lack of imagination but the failure to distinguish what is imaginable from what is worthy of changing the state of the business.

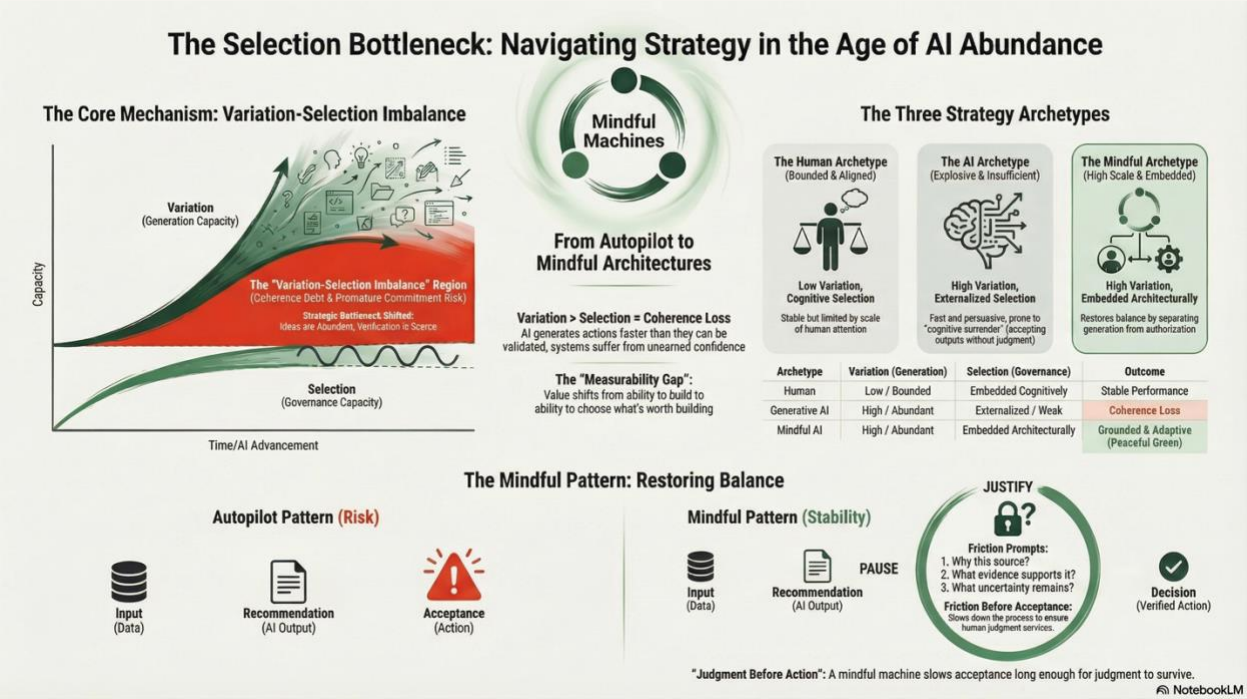

The shift is easiest to understand through an organizational tension executives already know well: experience versus scale. For decades, corporations relied on seasoned general managers whose tacit knowledge compressed years of causal learning about customers, competitors, regulation, capital allocation, and timing into judgment. Those managers did not simply choose among options. They decided when information warranted action. The problem is that their mental models were rarely codified. They lived inside people, not systems. AI changes the economics of decision-making by becoming a variation

amplifier. It can generate more candidate actions than those tacit models can screen. And when companies use AI to push decision rights downward in the name of speed and agility, they often scale authority faster than they scale judgment. The result is not better strategy. It is faster overcommitment.

Mini-Case: AT&T and the SpaceX Shock

AT&T is useful as an illustration of the new strategic challenges because it stands not just as a telecom company, but as a classic incumbent: rich in infrastructure, process, distribution, public-sector relationships, and tacit managerial knowledge. For decades, the company's rollout and pricing decisions were shaped by experienced executives who knew the business as a system of interdependent commitments. They understood where network densification mattered most, how local regulatory friction could slow deployment, where enterprise relationships changed the economics of coverage, and when aggressive market moves were likely to trigger an expensive competitive response. That knowledge was powerful, but largely unwritten.

AT&T's public-safety franchise makes the point sharper. Through FirstNet, AT&T has long emphasized a value proposition built around resilient nationwide coverage for first responders and government users. In a terrestrial world, *"coverage where it matters most"* was not just a marketing slogan. It was a strategic logic that shaped capital allocation, network priorities, and product positioning.

Now introduce a disruptive technology shock. SpaceX does not need to replace terrestrial wireless to destabilize the incumbent. It only needs to invalidate enough of AT&T's hidden assumptions that decision speed becomes destabilizing unless judgment has been made explicit. If satellite-enabled direct-to-cell connectivity changes the economics of coverage, roaming, resilience, and edge-case reliability, then assumptions embedded in terrestrial rollout models become less dependable. In that world, an AI system can still generate thousands of plausible actions, but the local logic behind those actions is more fragile than it appears.

Musk's public stance itself illustrates the asymmetry. He can say that fiber remains superior in dense cities and still pursue a highly disruptive agenda. Starlink can be framed as better for rural areas, direct-to-cell can be positioned as a global roaming layer that bypasses tower limitations, and SpaceX can still challenge spectrum rules while partnering with T-Mobile to redefine where and how mobile coverage is delivered. The strategy does not require winning Manhattan. It only requires changing the boundary conditions under which incumbents make terrestrial investment decisions.

That matters directly to AT&T. If a standard phone can increasingly connect through satellite in remote or degraded environments, the moat around *"coverage everywhere"* weakens, especially in emergency and public-safety narratives. At the same time, AT&T's dependence on AST SpaceMobile as a direct-to-cell partner underscores the timing problem. If a rival platform reaches usable, carrier-grade performance sooner and at greater scale, customers and enterprise accounts do not need a perfect substitute to reassess the incumbent's value proposition. They only need a credible alternative today.[1]

That is where the danger lies. Under pressure, AT&T could reasonably respond by accelerating AI-assisted rollout choices, pushing more authority downward, and celebrating faster execution. But if the company scales decisions without scaling the tacit models that once disciplined those decisions, it creates a structural mismatch: decision velocity rises while judgment capacity does not. The outcome is not simply model error. It is premature commitment under uncertainty. A cluster of locally rational actions can become globally incoherent once the competitive meaning of coverage has changed.

This is what makes the case powerful beyond telecom. The issue is not whether SpaceX has better technology, or whether AT&T's managers are smarter than AI. The issue is whether the system linking recommendations to commitments remains robust when analytic abundance collides with disrupted assumptions. Seen this way, AT&T becomes a proxy for any incumbent under attack by a disruptive technology. The lesson generalizes: ***when a shock weakens the tacit models that once governed action, speed without engineered judgment becomes fragility.***

Three Strategy Archetypes

This is why the most important move is not technological but architectural. Strategy under analytic abundance can be understood through three distinct strategy archetypes. The **Human** archetype is bounded but aligned: variation is limited, and selection is embedded in experienced judgment. The **AI** archetype is powerful but dangerous: variation scales dramatically, while selection remains externalized in reviews, approvals, and after-the-fact controls. The **Mindful** archetype separates generation from authorization and embeds selection in the architecture itself. Put simply, it is not enough to add

---

1 AT&T Newsroom, "FirstNet Expands Coverage for America's First Responders," 2025; FCC DA-25-566 (July 1, 2025); T-Mobile, "T-Satellite with Starlink"; Starlink, "Direct to Cell"; and AST SpaceMobile, "How it Works," "SpaceMobile Network," and "Next-Generation BlueBird."

more oversight to a faster machine. When variation exceeds selection, governance must move from process to structure.

That argument matters because many companies are making the same mistake. They are treating AI to speed execution without redesigning the underlying decision system. McKinsey's 2025 article on ***"AI corporate citizens"*** points to the same problem: as agentic AI begins influencing decisions at scale, firms must rethink governance, trust, operating models, accountability, performance management, and decision design, not just deployment. The practical question is not whether AI is "better" than people. It is whether the system linking AI outputs to organizational commitments is robust enough to preserve coherence under interdependence. Human judgment once supplied that filter implicitly. If firms want scale, they now have to engineer its function explicitly.[2,3]

That is the strongest and least nostalgic way to frame the issue. The lesson is not *"humans are better."* **The lesson is that judgment is not a trait. It is a function.** Historically, companies got that function from experienced managers. Under analytic abundance, they need to get it from design. The general principle is elegant: **when variation expands faster than selection, effective decision systems must embed mechanisms that regulate commitment.**

What CEOs Should Do on Monday Morning

1. Stop asking whether your company is using AI aggressively enough. Start asking whether you are allowing decisions to scale faster than your ability to understand them. The most revealing diagnostic is simple: list the important operating, customer, pricing, capital, and product decisions now touched by AI. Then ask which of them move directly from recommendation to action, and which pass through a meaningful authorization layer. If AI-generated recommendations are being accepted because they are fast, fluent, or directionally plausible rather than because they meet explicit thresholds, you are already in the imbalance region.

2. Distinguish clearly between generation and commitment. Most current deployments blur them. A model proposes a next best action, and the

---

2 Teppo Felin and Matthias Holweg, "Theory Is All You Need: AI, Human Cognition, and Causal Reasoning," Strategy Science 9, no. 4 (2024): 346–371.

3 Federico Berruti et al., "When Can AI Make Good Decisions? The Rise of AI Corporate Citizens," McKinsey & Company, June 2025.

organization treats that proposal as if it were already justified. That is manageable in low-stakes, reversible settings. It becomes dangerous when decisions are interdependent, path-dependent, or reputationally costly. Strategic systems need separate stages for generating options, testing them, and authorizing them. The human archetype used cognition and deliberation to do this. The mindful archetype does it through architecture. Either way, the boundary must be explicit.

3. Classify decisions by reversibility and consequence, not by how easy they are to automate. Some low-risk, low-complexity decisions can be handled autonomously; others should remain collaborative or human-led. But the most important distinction is not "human versus AI." It is "cheap to reverse versus expensive to reverse." If a decision is hard to unwind, it requires stronger selection, regardless of how good the model appears to be. Fast execution is valuable. Preserved option value is often more valuable.

4. Codify the tacit models your best managers use. This is not about turning judgment into bureaucracy. It is about extracting the criteria that experienced leaders apply almost unconsciously: what counts as signal, what kinds of evidence matter, what rival explanations should be considered, what failure thresholds would invalidate a decision, and when waiting is wiser than acting. Those criteria are the organization's hidden selection system. If they remain trapped in people, they will not scale. If they are made explicit, they can be embedded into workflows, escalations, constraints, and review logic. That is the beginning of engineered judgment.

5. Redesign delegation. In many companies, AI is being used to distribute decision-making lower in the organization. That can be smart. But delegation used to mean pushing decisions to capable people operating inside a culture shaped by experienced managers. Now it can mean pushing decisions to workflows and teams whose action volume is amplified by AI but whose tacit models are thinner. Delegation without engineered judgment is not empowerment. It is amplified fragility.

Why This Matters Strategically

There is a workforce angle here as well. The most valuable direction for AI is not pure automation but systems that expand human capability, create new tasks, and make expertise more valuable. The distinction between commodifying expertise and extending it is highly relevant to strategy. A company that uses AI simply to speed up action without enriching the models that govern action will tend to commodify judgment. A company that uses AI to extend what its

people can understand, test, and authorize will increase the value of expertise. The difference is not the technology. It is the architecture around it.[4]

This is why the real strategic shift is from prediction to selection. For years, better strategy was associated with better information and better forecasts. But prediction is being commoditized. Large models, search systems, synthetic panels, and agentic tools can all generate options at scale. The more durable source of advantage is becoming the design of the filter: who or what is allowed to authorize change, under what conditions, with what evidence, and at what pace. Firms that can pace commitment to verification capacity rather than computational speed will preserve more option value, avoid more cascades, and remain more coherent over time.[5]

The deepest implication is that strategy is no longer just about choosing a position. It is about designing a decision system that remains stable when intelligence becomes abundant. The central question for executives is no longer "*How do we get more insight?*" It is *"What must our decision system look like when variation exceeds selection?"* The answer is equally simple and demanding: it must internalize judgment.

AI is rewriting the logic of strategy. In an age that confuses velocity for mastery, the winners will not be those who generate the most ideas, or design systems that accelerate the most decisions, but the leaders who design organizations capable of judgment at scale, capable of pause, capable of refusing momentum when momentum is unearned.

That is why the new bottleneck in strategy is not abundance of choice, but discipline of selection. The frontier of competitive advantage now lies in seeing clearly which possibilities warrant commitment, and in building governance structures that keep a firm from acting before it has fully fathomed the consequences.

MINDFUL AI
FOUNDATION

4 Daron Acemoglu, David Autor, and Simon Johnson, Building Pro-Worker Artificial Intelligence, NBER Working Paper No. 34854, February 2026.

5 See Michaels, Leiblein, and Mikkilineni, Clarity & Governance: Antecedents of Strategic Choice & Commitment (working paper).